arXiv:1402.4746v1 [cs.LG] 19 Feb 2014

Near-optimal-sample estimators for spherical Gaussian mixtures

Jayadev Acharya[*], Ashkan Jafarpour[†], Alon Orlitksy[‡], and Ananda Theertha Suresh[§]

University of California, San Diego

February 20, 2014

Abstract

Statistical and machine-learning algorithms are frequently applied to high-dimensional data. In many of these applications data is scarce, and often much more costly than computation time. We provide the first sample-efficient polynomial-time estimator for high-dimensional spherical Gaussian mixtures.

For mixtures of any k d-dimensional spherical Gaussians, we derive an intuitive spectral-estimator that uses $\mathcal{O}_k\big(\frac{d\log^2 d}{\epsilon^4}\big)$ samples and runs in time $\mathcal{O}_{k,\epsilon}(d^3\log^5 d)$, both significantly lower than previously known. The constant factor $\mathcal{O}_k$ is polynomial for sample complexity and is exponential for the time complexity, again much smaller than what was previously known. We also show that $\Omega_k\big(\frac{d}{\epsilon^2}\big)$ samples are needed for any algorithm. Hence the sample complexity is near-optimal in the number of dimensions.

We also derive a simple estimator for k-component one-dimensional mixtures that uses $\mathcal{O}\big(\frac{k\log\frac{k}{\epsilon}}{\epsilon^2}\big)$ samples and runs in time $\widetilde{\mathcal{O}}\Big(\big(\frac{k}{\epsilon}\big)^{3k+1}\Big)$. Our other technical contributions include a faster algorithm for choosing a density estimate from a set of distributions, that minimizes the ℓ_1 distance to an unknown underlying distribution.

[*]jacharya@ucsd.edu
[†]ashkan@ucsd.edu
[‡]alon@ucsd.edu
[§]asuresh@ucsd.edu

1 Introduction

1.1 Background

Meaningful information often resides in high-dimensional spaces: voice signals are expressed in many frequency bands, credit ratings are influenced by multiple parameters, and document topics are manifested in the prevalence of numerous words. Some applications, such as topic modeling and genomic analysis consider data in over 1000 dimensions, [17, 44].

Typically, information can be generated by different types of sources: voice is spoken by men or women, credit parameters correspond to wealthy or poor individuals, and documents address topics such as sports or politics. In such cases the overall data follow a mixture distribution [26, 36, 38].

Mixtures of high-dimensional distributions are therefore central to the understanding and processing of many natural phenomena. Methods for recovering the mixture components from the data have consequently been extensively studied by statisticians, engineers, and computer scientists.

Initially, heuristic methods such as expectation-maximization (EM) were developed [27, 35]. Over the past decade, more rigorous algorithms were derived to recover mixtures of d-dimensional spherical Gaussians [5, 7, 11, 12, 21, 42], general Gaussians [2, 4, 6, 10, 22, 29], and other log-concave distributions [23]. Many of these algorithms consider mixtures where the ℓ_1 distance between the mixture components is $2 - o_d(1)$, namely approaches the maximum of 2 as d increases. They identify the distribution components in time and samples that grow polynomially in the dimension d. Recently, [22, 29] showed that any d-dimensional Gaussian mixture can be recovered in polynomial time. However, their algorithm uses $> d^{100}$ time and samples.

A different approach that avoids the large component-distance requirement and the high time and sample complexity, considers a slightly more relaxed notion of approximation, sometimes called *PAC learning*. PAC learning [24] does not approximate each mixture component, but instead derives a mixture distribution that is close to the original one. Specifically, given a distance bound $\epsilon > 0$, error probability $\delta > 0$, and samples from the underlying mixture $\mathbf{f}$, where we use boldface letters for d-dimensional objects, PAC learning seeks a mixture estimate $\hat{\mathbf{f}}$ with at most k components such that $D(\mathbf{f}, \hat{\mathbf{f}}) \le \epsilon$ with probability $\ge 1 - \delta$, where $D(\cdot, \cdot)$ is some given distance measure, for example ℓ_1 distance or KL divergence. This notion of estimation is also known as *proper learning* in the literature.

An important and extensively studied special case of mixture distributions are *spherical-Gaussians* [5, 7, 11, 12, 21, 42], where different coordinates have the same variance, though potentially different means. Due to their simple structure, they are easier to analyze and under a minimum-separation assumption have provably-practical algorithms for clustering and parameter estimation [7, 11, 12, 42].

1.2 Sample complexity

Reducing the number of samples is of great practical significance. For example, in topic modeling every sample is a whole document, in credit analysis every sample is a person's credit history, and in genetics, every sample is a human DNA. Hence samples can be very scarce and obtaining them can be very costly. By contrast, current CPUs run at several Giga Hertz, hence samples are typically much more scarce of a resource than time.

Note that for one-dimensional statistical problems, the need for sample-efficient algorithms has been broadly recognized. The sample complexity of many problems is known quite accurately, often to within a constant factor. For example, for discrete distributions over $\{1, \ldots, s\}$, an approach proposed in [32] and its modifications were used in [40, 41] to estimate the probability multiset using $\Theta(s/\log s)$ samples.

Learning one-dimensional m-modal distributions over $\{1,\ldots,s\}$ requires $\Theta(m\log(s/m)/\epsilon^3)$ samples [14]. Similarly, one-dimensional mixtures of k structured distributions (log-concave, monotone hazard rate, and unimodal) over $\{1,\ldots,s\}$ can be learned with $\mathcal{O}(k/\epsilon^4)$, $\mathcal{O}(k\log(s/\epsilon)/\epsilon^4)$, and $\mathcal{O}(k\log(s)/\epsilon^4)$ samples, respectively, and these bounds are tight up to a factor of ϵ [31].

Compared to one dimensional problems, in high dimensions there is a polynomial gap in the sample complexity. For example, for learning spherical Gaussian mixtures, the number of samples required by previous algorithms is $\mathcal{O}(d^{12})$ for $k = 2$ components, and increased exponentially with k [19]. In this paper we bridge this gap, by constructing near-linear sample complexity estimators.

1.3 Previous and new results

Our main contribution is PAC learning d dimensional Gaussian mixtures with near-linear samples. We show few auxiliary results for one-dimensional Gaussians.

1.3.1 d-dimensional Gaussian mixtures

Several papers considered PAC learning of discrete- and Gaussian-product mixtures. [20] considered mixtures of two d-dimensional Bernoulli products where all probabilities are bounded away from 0. They showed that this class is PAC learnable in $\widetilde{\mathcal{O}}(d^2/\epsilon^4)$ time and samples, where the $\widetilde{\mathcal{O}}$ notation hides logarithmic factors. [18] eliminated the probability constraints and generalized the results from binary to arbitrary discrete alphabets, and from 2 to k mixture components. They showed that mixtures of k discrete products are PAC learnable in $\widetilde{\mathcal{O}}\big((d/\epsilon)^{2k^2(k+1)}\big)$ time, and although they did not explicitly mention sample complexity, their algorithm uses $\widetilde{\mathcal{O}}\big((d/\epsilon)^{4(k+1)}\big)$ samples. [19] generalized these results to Gaussian products, showing in particular that mixtures of k Gaussians, where the difference between the means normalized by the ratio of standard deviations is bounded by B, are PAC learnable in $\widetilde{\mathcal{O}}\big((dB/\epsilon)^{2k^2(k+1)}\big)$ time, and can be shown to use $\widetilde{\mathcal{O}}\big((dB/\epsilon)^{4(k+1)}\big)$ samples. These algorithms consider the KL divergence between the distribution and its estimate, but it can be shown that the ℓ_1 distance would result in similar complexities. It can also be shown that these algorithms or their simple modifications have similar time and sample complexities for spherical Gaussians as well.

Our main contribution shows that mixtures of spherical-Gaussians are PAC learnable in ℓ_1 distance with sample complexity that is nearly linear in the dimension. Specifically, Theorem 8 shows that mixtures of k spherical-Gaussian distributions can be learned in

$$n = \mathcal{O}\left(\frac{dk^9}{\epsilon^4}\log^2\frac{d}{\delta}\right) = \mathcal{O}_{k,\epsilon}(d\log^2 d)$$

samples and

$$\mathcal{O}\Big(n^2 d\log n + d^2\Big(\frac{k^7}{\epsilon^3}\log\frac{d}{\delta}\Big)^{k^2}\Big) = \widetilde{\mathcal{O}}_{k,\epsilon}(d^3).$$

time. Observe that recent algorithms typically construct the covariance matrix [19, 42], hence require $\geq nd^2$ time. In that sense, for small values of k, the time complexity we derive is comparable to the best such algorithms can hope for. Observe also that the exponential dependence on k is of the form $d^2\Big(\frac{k^7}{\epsilon^3}\log\frac{d}{\delta}\Big)^{k^2}$, which is significantly lower than the $d^{\mathcal{O}(k^3)}$ dependence in previous results.

By contrast, Theorem 2 shows that PAC learning k-component spherical Gaussian mixtures require $\Omega(dk/\epsilon^2)$ samples for any algorithm, hence our distribution learning algorithms are nearly sample optimal. In addition, their time complexity significantly improves on previously known ones.

1.3.2 One-dimensional Gaussian mixtures

Independently and around the same time as this work [15] showed that mixtures of two one-dimensional Gaussians can be learnt with $\widetilde{\mathcal{O}}(\epsilon^{-2})$ samples and in time $\mathcal{O}(\epsilon^{-7.01})$. We provide a natural estimator for learning mixtures of k one dimensional Gaussians using some basic properties of Gaussian distributions and show that mixture of any k-one dimensional Gaussians can be learnt with $\widetilde{\mathcal{O}}(k\epsilon^{-2})$ samples and in time $\widetilde{\mathcal{O}}\left(\left(\frac{k}{\epsilon}\right)^{3k+1}\right)$.

1.4 The approach and technical contributions

The popular SCHEFFE estimator takes a collection $\mathcal{F}$ of distributions and uses $\mathcal{O}(\log|\mathcal{F}|)$ independent samples from an underlying distribution $\mathbf{f}$ to find a distribution in $\mathcal{F}$ whose distance from $\mathbf{f}$ is at most a constant factor larger than that of the distribution in $\mathcal{F}$ that is closet to f [16]. In Lemma 1, we lower the time complexity of the Scheffe algorithm from $\mathcal{O}(|\mathcal{F}|^2)$ time to $\widetilde{\mathcal{O}}(|\mathcal{F}|)$, helping us reduce the time complexity of our algorithms.

Our goal is therefore to construct a small class of distributions that is ϵ-close to any possible underlying distribution. For simplicity, consider spherical Gaussians with the same variance and means bounded by B. Take the collection of all distributions derived by quantizing the means of all components in all coordinates to ϵ_m accuracy, and quantizing the weights to ϵ_w accuracy. It can be shown that to get distance ϵ from the underlying distribution, it suffices to take $\epsilon_m, \epsilon_w \leq 1/\mathrm{poly}_\epsilon(dk)$. There are at most $\left(\frac{B}{\epsilon_m}\right)^{dk} \cdot \left(\frac{1}{\epsilon_w}\right)^{k} = 2^{\widetilde{\mathcal{O}}_\epsilon(dk)}$ possible combinations of the k mean vectors and weights. Hence SCHEFFE implies an exponential-time algorithm with sample complexity $\widetilde{\mathcal{O}}(dk)$.

To reduce the dependence on d, one can approximate the span of the k mean vectors. This reduces the problem from d to k dimensions, allowing us to consider a distribution collection of size $2^{\mathcal{O}(k^2)}$, with SCHEFFE sample complexity of just $\mathcal{O}(k^2)$. [18, 19] constructs the sample correlation matrix and uses k of its columns to approximate the span of mean vectors. This approach requires the k columns of the sample correlation matrix to be very close to the actual correlation matrix, and thus requires a lot more samples.

We derive a spectral algorithm that uses the top k eigenvectors of the sample covariance matrix to approximate the span of the k mean vectors. Since we use the entire covariance matrix instead of just k columns, a weaker concentration is sufficient and we gain on the sample complexity.

Using recent tools from non-asymptotic random matrix theory [3, 39, 43], we show that the approximation of the span of the means converges in $\widetilde{\mathcal{O}}(d)$ samples. This result allows us to address most "reasonable" distributions, but still there are some "corner cases" that need to be analyzed separately. To address them, we modify some known clustering algorithms such as single-linkage, and spectral projections. While the basic algorithms were known before, our contribution here, which takes a fair bit of effort and space, is to show that judicious modifications of the algorithms and rigorous statistical analysis yield polynomial time algorithms with near optimal sample complexity.

Our approach applies most directly to mixtures of spherical Gaussians. We provide a simple and practical recursive clustering and spectral algorithm that estimates all such distributions in $\mathcal{O}_k(d\log^2 d)$ samples.

The paper is organized as follows. In Section 2, we introduce notations, describe results on the Scheffe estimator, and state a lower bound. In Section 3, we present the algorithm for k-spherical Gaussians. In Section 4 we show a simple learning algorithm for one-dimensional Gaussian mixtures. To preserve readability, most of the technical details and proofs are given in the appendix.

2 Preliminaries

2.1 Notation

For arbitrary product distributions $\mathbf{p}_1, \ldots, \mathbf{p}_k$ over a d dimensional space let $p_{j,i}$ be the distribution of $\mathbf{p}_j$ over coordinate i, and let $\mu_{j,i}$ and $\sigma_{j,i}$ be the mean and variance of $p_{j,i}$ respectively. Let $\mathbf{f} = (w_1, \ldots, w_k, \mathbf{p}_1, \ldots, \mathbf{p}_k)$ be the mixture of these distributions with mixing weights $w_1, \ldots, w_k$. We denote estimates of a quantity $\mathbf{x}$ by $\hat{\mathbf{x}}$. It can be empirical mean or a more complex estimate. $||\cdot||$ denotes the spectral norm of a matrix and $||\cdot||_2$ denotes the ℓ_2 norm of a vector.

2.2 Selection from a pool of distributions

Many algorithms for learning mixtures over the domain $\mathcal{X}$ first obtain a *small* collection of mixtures distributions $\mathcal{F}$ and then perform Maximum Likelihood test using the samples to output a distribution [14,18,20]. Our algorithm also obtains a set of distributions containing at least one that is close to the underlying in ℓ_1 distance. The estimation problem now reduces to the following. Given a class $\mathcal{F}$ of distributions and samples from an unknown distribution $\mathbf{f}$, find a distribution in $\mathcal{F}$ that is *close* to $\mathbf{f}$. Let $D(\mathbf{f}, \mathcal{F}) \stackrel{\text{def}}{=} \min_{\mathbf{f}_i \in \mathcal{F}} D(\mathbf{f}, \mathbf{f}_i)$.

The well-known Scheffe's method [16] uses $\mathcal{O}(\epsilon^{-2} \log |\mathcal{F}|)$ samples from the underlying distribution $\mathbf{f}$, and in time $\mathcal{O}(\epsilon^{-2}|\mathcal{F}|^2 T \log |\mathcal{F}|)$ outputs a distribution in $\mathcal{F}$ with ℓ_1 distance of at most $9.1 \max(D(\mathbf{f}, \mathcal{F}), \epsilon)$ from $\mathbf{f}$, where T is the time required to compute the probability of an $x \in \mathcal{X}$ by a distribution in $\mathcal{F}$. A naive application of this algorithm requires time quadratic in the number of distributions in $\mathcal{F}$. We propose a variant of this, that works in near linear time, albeit requiring slightly more samples. More precisely,

Lemma 1 (Appendix B)**.** *Let $\epsilon > 0$. For some constant c, given $\frac{c}{\epsilon^2} \log\big(\frac{|\mathcal{F}|}{\delta}\big)$ independent samples from a distribution $\mathbf{f}$, with probability $\geq 1 - \delta$, the output $\hat{\mathbf{f}}$ of* MODIFIED SCHEFFE *$D(\hat{\mathbf{f}}, \mathbf{f}) \leq 1000 \max(\epsilon, D(\mathbf{f}, \mathcal{F}))$. Furthermore, the algorithm runs in time $\mathcal{O}\big(\frac{|\mathcal{F}| T \log(|\mathcal{F}|/\delta)}{\epsilon^2}\big)$.*

We therefore find a *small* class $\mathcal{F}$ with at least one distribution close to the underlying mixture. For our problem of estimating k component mixtures in d-dimensions, $T = \mathcal{O}(dk)$ and $|\mathcal{F}| = \widetilde{\mathcal{O}}_{k,\epsilon}(d^2)$. Note that we have not optimized the constant 1000 in the above lemma.

2.3 Lower bound

Using Fano's inequality, we show an information theoretic lower bound of $\Omega(dk/\epsilon^2)$ samples to learn k-component d-dimensional mixtures of spherical Gaussians for any algorithm. More precisely,

Theorem 2 (Appendix C)**.** *Any algorithm that learns all k-component d-dimensional spherical Gaussian mixtures up to ℓ_1 distance ϵ with probability $\geq 1/2$ requires at least $\Omega(\frac{dk}{\epsilon^2})$ samples.*

3 Mixtures in d dimensions

3.1 Description of LEARN k-SPHERE

Algorithm LEARN K-SPHERE learns mixtures of k spherical Gaussians using near-linear samples. For clarity, we assume that all components have the same variance σ^2, *i.e.*, $\mathbf{p}_i = N(\boldsymbol{\mu}_i, \sigma^2 \mathbb{I}_d)$ for $1 \leq i \leq k$. A modification of this algorithm works for components with different variances. The core ideas are same and we include it in the final version of the paper.

The easy part of the algorithm is estimating σ^2. If $\mathbf{X}(1)$ and $\mathbf{X}(2)$ are two samples from the same component, then $\mathbf{X}(1) - \mathbf{X}(2)$ is distributed $N(0, 2\sigma^2\mathbb{I}_d)$. Hence for large d, $\|\mathbf{X}(1) - \mathbf{X}(2)\|_2^2$ concentrates around $2d\sigma^2$. By the pigeon-hole principle, given $k+1$ samples, two of them are from the same component. Therefore, the minimum pairwise distance between $k+1$ samples is close to $2d\sigma^2$. This constitutes the first step of our algorithm.

We now concentrate on estimating the means. As stated in the introduction, given the span of the mean vectors $\boldsymbol{\mu}_i$, we can grid the k dimensional span to the required accuracy ϵ_g and use SCHEFFE, to obtain a polynomial time algorithm. One of the natural and well-used methods to estimate the span of mean vectors is using the correlation matrix [42]. Consider the correlation-type matrix,

$$S = \frac{1}{n}\sum_{i=1}^{n} \mathbf{X}(i)\mathbf{X}(i)^t - \sigma^2\mathbb{I}_d.$$

In expectation, the fraction of terms from $\mathbf{p}_i$ is w_i. Furthermore for a sample X from a particular component j,

$$\mathbb{E}[\mathbf{X}\mathbf{X}^t] = \sigma^2\mathbb{I}_d + \boldsymbol{\mu}_j\boldsymbol{\mu}_j{}^t.$$

It follows that

$$\mathbb{E}[S] = \sum_{j=1}^{k} w_j\boldsymbol{\mu}_j\boldsymbol{\mu}_j{}^t.$$

Therefore, as $n \to \infty$, the matrix S converges to $\sum_{j=1}^{k} w_j\boldsymbol{\mu}_j\boldsymbol{\mu}_j{}^t$, and its top k eigenvectors span of means.

While the above intuition is well understood, the number of samples necessary for convergence is not well studied. Ideally, irrespective of the values of the means, we wish $\widetilde{\mathcal{O}}(d)$ samples to be sufficient for the convergence. However this is not true, as we demonstrate by a simple example.

Example 3. *Consider the special case, $d = 1$, $k = 2$, $\sigma^2 = 1$, $w_1 = w_2 = 1/2$, and the difference of means $|\mu_1 - \mu_2| = L$ for a large $L \gg 1$. Given this prior information, one can estimate the the average of the mixture, that yields $\frac{\mu_1+\mu_2}{2}$. Solving equations obtained by $\mu_1 + \mu_2$ and $\mu_1 - \mu_2 = L$, yields μ_1 and μ_2. The variance of the mixture is $1 + \frac{L^2}{4} > \frac{L^2}{4}$. With additional Chernoff type bounds, one can show that given n samples the error in estimating the average is*

$$|\mu_1 + \mu_2 - \hat{\mu}_1 - \hat{\mu}_2| \approx \Theta\left(\frac{L}{\sqrt{n}}\right).$$

Therefore to estimate the means to a small accuracy we need $n \geq L^2$, i.e., more the separation, more samples are necessary.

A similar phenomenon happens in the convergence of the correlation matrices, where the variances of quantities of interest increases with separation. In other words, for the span to be accurate the number of samples necessary increases with the separation. To overcome this phenomenon, a natural idea is to cluster the Gaussians such that the means of components in the same cluster are close and then apply SCHEFFE on the span within each cluster.

Even though spectral clustering algorithms are studied in [2,42], they assume that the weights are strictly bounded away from 0, which does not hold here. We use a simple recursive clustering algorithm that takes a cluster C with average $\overline{\boldsymbol{\mu}}(C)$. If there is a component in the cluster such that $\sqrt{w_i}\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\|_2$ is $\Omega(\log(n/\delta))$, then the algorithm divides the cluster into two nonempty clusters without any mis-clustering.

For technical reasons similar to the above example, we also use a coarse clustering algorithm that ensures that the mean separation is $\widetilde{\mathcal{O}}(d^{1/4})$ within each cluster. The algorithm can be summarized as:

1. **Variance estimation:** Use first $k+1$ samples and estimate the minimum distance among sample-pairs to estimate σ^2.

2. **Coarse clustering:** Using a single-linkage algorithm, group the samples such that within each cluster formed, the mean separation is smaller than $\widetilde{\mathcal{O}}(d^{1/4})$.

3. **Recursive clustering**: As long as there is a cluster that has samples from more than one component with means *far* apart, (described by a condition on the norm of its covariance matrix in the algorithm) estimate its largest eigenvector and project samples of this cluster onto this eigenvector and cluster them. This hierarchical method is continued until there are clusters that contain close-by-components.

4. **Search in the span:** The resulting clusters contain components that are close-by, *i.e.,* $\left\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\right\|_2 < \mathcal{O}(k^{3/2}\hat{\sigma}^2 \log \frac{n}{\delta})$. We approximate the span of means by the top $k-1$ eigenvectors and the mean vector, and perform an exhaustive search using MODIFIED SCHEFFE.

We now describe these steps stating the performance of each step.

Algorithm LEARN K-SPHERE
Input: n samples $\mathbf{x}(1), \mathbf{x}(2), \ldots, \mathbf{x}(n)$ from $\mathbf{f}$ and ϵ.

1. **Sample variance:** $\hat{\sigma}^2 = \min_{a \neq b : a,b \in [k+1]} ||\mathbf{x}(a) - \mathbf{x}(b)||_2^2 / 2d$.

2. **Coarse single-linkage clustering:** Start with each sample as a cluster,

   - While $\exists$ two clusters with squared-distance $\leq 2d\hat{\sigma}^2 + 23\hat{\sigma}^2\sqrt{d \log \frac{n^2}{\delta}}$, merge them.

3. **Recursive spectral-clustering:** While there is a new cluster C with $|C| \geq n\epsilon/5k$ and spectral norm of its sample covariance matrix $\geq 12k^2\hat{\sigma}^2 \log n^3/\delta$,

   - Use $n\epsilon/8k^2$ of the samples to find the largest eigenvector and discard these samples.
   - Project the remaining samples on the largest eigenvector.
   - Perform single-linkage in the projected space (as before) till the distance between clusters $> 3\hat{\sigma}\sqrt{\log n^2 k/\delta}$ creating new clusters.

4. **Exhaustive search**: Let $\epsilon_g = \epsilon/(16k^{3/2})$, $L = 200\sqrt{k^4\epsilon^{-1} \log \frac{n^2}{\delta}}$, and $G = \{-L, \ldots, -\epsilon_g, 0, \epsilon_g, 2\epsilon_g, \ldots L\}$. Let $W = \{0, \epsilon/(4k), 2\epsilon/(4k), \ldots 1\}$ and $\Sigma \stackrel{\text{def}}{=} \{\sigma^2 : \sigma^2 = \hat{\sigma}^2(1 + i/d) \forall -d < i \leq d\}$.

   - For each cluster C find its top $k-1$ eigenvectors $\mathbf{u}_1, \mathbf{u}_2 \ldots \mathbf{u}_{k-1}$ and let $\text{Span}(C) = \{\hat{\overline{\boldsymbol{\mu}}}(C) + \sum_{i=1}^{k-1} g_i \hat{\sigma} \mathbf{u}_i : g_1, g_2 \ldots g_{k-1} \in G\}$.
   - Let $\text{Span} = \{\text{Span}(C) : |C| \geq n\epsilon/5k\}$.
   - For all $w_i' \in W$, $\sigma'^2 \in \Sigma$, $\hat{\boldsymbol{\mu}}_i \in \text{Span}$, add $\{(w_1', \ldots, w_{k-1}', 1 - \sum_{i=1}^{k-1} w_i', N(\hat{\boldsymbol{\mu}}_1, \sigma'^2), \ldots, N(\hat{\boldsymbol{\mu}}_k, \sigma'^2)\}$ in $\mathcal{F}$.

5. Run MODIFIED SCHEFFE on $\mathcal{F}$ and output the resulting distribution.

3.2 Sketch of correctness

To simplify the bounds and expressions, we assume that $d > 1000$ and $\delta \geq \min(2n^2e^{-d/10}, 1/3)$. For smaller values of δ, we run the algorithm with error $1/3$ and repeat it $\mathcal{O}(\log \frac{1}{\delta})$ times to choose a set of candidate mixtures $\mathcal{F}_\delta$. By Chernoff-bound with error $\leq \delta$, $\mathcal{F}_\delta$ contains a mixture ϵ-close to $\mathbf{f}$. Finally, we run MODIFIED SCHEFFE on $\mathcal{F}_\delta$ to obtain a mixture that is close to $\mathbf{f}$. By the union bound and Lemma 1, the error is $\leq 2\delta$.

Variance estimation: Let $\hat{\sigma}$ be the variance estimate from step 1. In high dimensions, the difference between two random samples from a Gaussian concentrates. This is made precise in the next lemma which states $\hat{\sigma}$ is a good estimate of the variance. Then the following is a simple application of Gaussian tail bounds.

Lemma 4 (Appendix D.1)**.** *Given n samples from the k-component mixture, with probability $1 - 2\delta$,*

$$|\hat{\sigma}^2 - \sigma^2| \leq 2.5\sigma^2\sqrt{\frac{\log(n^2/\delta)}{d}}.$$

Coarse single-linkage clustering: The second step is a single-linkage routine that clusters mixture components with *far* means. Single-linkage is a simple clustering scheme that starts out with each data point as a cluster, and at each step merges the two that are closest to form larger clusters. The algorithm stops when the distance between clusters is larger than a pre-specified threshold.

Suppose the samples are generated by an one-dimensional mixture of k components that are *far*, then *with high probability*, when the algorithm generates k clusters and all the samples within a cluster are generated by a single component. More precisely, if $\forall i, j \in [k]$, $|\mu_i - \mu_j| = \Omega(\sigma \log n)$, then all the n samples concentrate around their respective means and the separation between any two samples from different components would be larger than the largest separation between any two samples from the same component. Hence for a suitable value of threshold, single-linkage correctly identifies the clusters. For d-dimensional Gaussian mixtures a similar notion holds true, with minimum separation $\Omega(d^{1/4} \log \frac{n}{\delta})$. More precisely,

Lemma 5 (Appendix D.2)**.** *After Step 2 of* LEARN K-SPHERE*, with probability $\geq 1 - 2\delta$, all samples from each component will be in the same cluster and the maximum distance between two components within each cluster is $\leq 10k\sigma\big(d \log \frac{n^2}{\delta}\big)^{1/4}$.*

Recursive spectral-clustering: The clusters formed at this step consists of components with mean separation $\mathcal{O}(d^{1/4} \log \frac{n}{\delta})$. We now recursively zoom into the clusters formed and show that it is possible to cluster the components with much smaller mean separation. Note that since the matrix is symmetric, the largest magnitude of the eigenvalue is same as the spectral norm. We first find the largest eigenvector of

$$S(C) \stackrel{\text{def}}{=} \frac{1}{|C|}\Big(\sum_{\mathbf{x} \in C} (\mathbf{x} - \hat{\overline{\boldsymbol{\mu}}}(C))(\mathbf{x} - \hat{\overline{\boldsymbol{\mu}}}(C))^t\Big) - \hat{\sigma}^2 \mathbb{I}_d,$$

which is the sample covariance matrix with its diagonal term reduced by $\hat{\sigma}^2$. If there are two components with means far apart, then using single-linkage we divide the cluster into two. The following lemma shows that this step performs accurate clustering of components with means well separated.

Lemma 6 (Appendix D.3)**.** *Let $n \geq c \cdot \frac{dk^4}{\epsilon} \log \frac{n^3}{\delta}$. After recursive clustering, with probability $\geq 1 - 4\delta$. the samples are divided into clusters such that for each component i within any cluster C, $\sqrt{w_i}\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\|_2 \leq 25\sigma\sqrt{k^3 \log \frac{n^3}{\delta}}$. Furthermore, all the samples from one component remain in a single cluster.*

Exhaustive search and Scheffe: After step 3, all clusters have a small weighted radius $\sqrt{w_i}\left\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\right\|_2 \leq 25\sigma\sqrt{k^3 \log \frac{n^3}{\delta}}$, the the eigenvectors give an accurate estimate of the span of $\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)$ within each cluster. More precisely,

Lemma 7 (Appendix D.4)**.** *Let $n \geq c \cdot \frac{dk^9}{\epsilon^4} \log^2 \frac{d}{\delta}$ for some constant c. After step 3, with probability $\geq 1 - 7\delta$ the following holds: if $|C| \geq n\epsilon/5k$, then the projection of $[\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)]/\left\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\right\|_2$ on the space orthogonal to the span of top $k-1$ eigenvectors has magnitude $\leq \frac{\epsilon\sigma}{8\sqrt{2}k\sqrt{w_i}\left\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\right\|_2}$.*

We now have accurate estimates of the spans of the clusters and each cluster has components with close means. It is now possible to grid the set of possibilities in each cluster to obtain a set of distributions such that one of them is close to the underlying. There is a trade-off between a dense grid to obtain a good estimation and the computation time required. The final step takes the sparsest grid possible to ensure an error $\leq \epsilon$. This is quantized below.

Theorem 8 (Appendix D.5)**.** *Let $n \geq c \cdot \frac{dk^9}{\epsilon^4} \log^2 \frac{d}{\delta}$ for some constant c. Then Algorithm* LEARN K-SPHERE *with probability $\geq 1 - 9\delta$, outputs a distribution $\hat{\mathbf{f}}$ such that $D(\hat{\mathbf{f}}, \mathbf{f}) \leq 1000\epsilon$. Furthermore, the algorithm runs in time $\mathcal{O}\Big(n^2 d \log n + d^2\Big(\frac{k^7}{\epsilon^3} \log \frac{d}{\delta}\Big)^{k^2}\Big)$.*

Note that the run time is calculated based on the efficient implementation of single-linkage [37] and the exponential term is not optimized. We now study mixtures in one-dimension and provide an estimator using MODIFIED SCHEFFE.

4 Mixtures in one dimension

Over the past decade estimating one dimensional distributions has gained significant attention [1, 13–15, 30, 31, 33, 41]. We now provide a simple estimator for learning one dimensional mixtures using the MODIFIED SCHEFFE estimator proposed earlier. The d-dimension estimator uses spectral projections to find the span of means, whereas for one dimension case, we use a simple observation on properties of samples from Gaussians for estimation. Formally, given samples from f, a mixture of Gaussian distributions $p_i \stackrel{\text{def}}{=} N(\mu_i, \sigma_i^2)$ with weights $w_1, w_2, \ldots w_k$, our goal is to find a mixture $\hat{f} = (\hat{w}_1, \hat{w}_2, \ldots \hat{w}_k, \hat{p}_1, \hat{p}_2, \ldots \hat{p}_k)$ such that $D(f, \hat{f}) \leq \epsilon$. Note that we make no assumption on the weights, means or the variances of the components.

We provide an algorithm that, using $\widetilde{\mathcal{O}}(k\epsilon^{-2})$ samples and in time $\widetilde{\mathcal{O}}(k\epsilon^{-3k-1})$, outputs an estimate that is at most ϵ from the underlying in ℓ_1 distance with probability $\geq 1 - \delta$. Our algorithm is an immediate consequence of the following observation for samples from a Gaussian distribution.

Lemma 9. *Given n independent samples $x_1, \ldots, x_n$ from $N(\mu, \sigma^2)$, there are two samples x_j, x_k such that $|x_j - \mu| \leq \sigma\frac{7\log 2/\delta}{2n}$ and $|x_j - x_k - \sigma| \leq 2\sigma\frac{7\log 2/\delta}{2n}$ with probability $\geq 1 - \delta$.*

Proof The density of $N(\mu, \sigma^2)$ is $\geq (7\sigma)^{-1}$ in the interval $[\mu - \sqrt{2}\sigma, \mu + \sqrt{2}\sigma]$. Therefore, the probability that a sample occurs in the interval $\mu - \epsilon\sigma, \mu + \epsilon\sigma$ is $\geq 2\epsilon/7$. Hence, the probability that none of the n samples occurs in $[\mu - \epsilon\sigma, \mu + \epsilon\sigma]$ is $\leq (1 - 2\epsilon/7)^n \leq e^{-2n\epsilon/7}$. If $\epsilon \geq \frac{7\log 2/\delta}{2n}$, then the probability that none of the samples occur in the interval is $\leq \delta/2$. A similar argument shows that there is a sample within interval, $[\mu + \sigma - \epsilon\sigma, \mu + \sigma + \epsilon\sigma]$, proving the lemma.

The above observation can be translated into selecting a pool of candidate distributions such that one of the distributions is close to the underlying distribution.

Lemma 10. *Given $n \geq \frac{120k \log \frac{4k}{\delta}}{\epsilon}$ samples from a mixture f of k Gaussians. Let $S = \{N(x_j, (x_j - x_k)^2) : 1 \leq j, k \leq n\}$ be a set of Gaussians and $W = \{0, \frac{\epsilon}{2k}, \frac{2\epsilon}{2k} \ldots, 1\}$ be the set of weights. Let*

$$\mathcal{F} \stackrel{\text{def}}{=} \{\hat{w}_1, \hat{w}_2, \ldots, \hat{w}_{k-1}, 1 - \sum_{i=1}^{k-1} \hat{w}_i, \hat{p}_1, \hat{p}_2, \ldots \hat{p}_k : \hat{w}_i \in W, \hat{p}_i \in S\}$$

be a set of $n^{2k}\left(\frac{2k}{\epsilon}\right)^{k-1} \leq n^{3k-1}$ candidate mixture distributions. There exists a $\hat{f} \in \mathcal{F}$ such that $D(f, \hat{f}) \leq \epsilon$.

Proof. Let $f = (w_1, w_2, \ldots w_k, p_1, p_2, \ldots p_k)$. For $\hat{f} = (\hat{w}_1, \hat{w}_2, \ldots, \hat{w}_{k-1}, 1 - \sum_{i=1}^{k-1} \hat{w}_i, \hat{p}_1, \hat{p}_2, \ldots \hat{p}_k)$, by the triangle inequality,

$$D(f, \hat{f}) \leq \sum_{i=1}^{k-1} 2|\hat{w}_i - w_i| + \sum_{i=1}^{k} w_i D(p_i, \hat{p}_i).$$

We show that there is a distribution in $\hat{f} \in \mathcal{F}$ such that the sum above is bounded by ϵ. Since we quantize the grids as multiples of $\epsilon/2k$, we consider distributions in $\mathcal{F}$ such that each $|\hat{w}_i - w_i| \leq \epsilon/4k$, and therefore $\sum_i |\hat{w}_i - w_i| \leq \frac{\epsilon}{2}$.

We now show that for each p_i there is a $\hat{p}_i$ such that $w_i D(p_i, \hat{p}_i) \leq \frac{\epsilon}{2k}$, thus proving that $D(f, \hat{f}) \leq \epsilon$. If $w_i \leq \frac{\epsilon}{4k}$, then $w_i D(p_i, \hat{p}_i) \leq \frac{\epsilon}{2k}$. Otherwise, let $w_i' > \frac{\epsilon}{4k}$ be the fraction of samples from p_i. By Lemma 9 and 14, with probability $\geq 1 - \delta/2k$,

$$\begin{aligned} D(p_i, \hat{p}_i)^2 &\leq 2\frac{(\mu_i - \mu_i')^2}{\sigma_i^2} + 16\frac{(\sigma_i - \sigma_i')^2}{\sigma_i^2} \\ &\leq \frac{25 \log^2 \frac{4k}{\delta}}{(nw_i')^2} + \frac{800 \log^2 \frac{4k}{\delta}}{(nw_i')^2} \\ &\leq \frac{825 \log^2 \frac{4k}{\delta}}{(nw_i')^2}. \end{aligned}$$

Therefore,

$$w_i D(p_i, \hat{p}_i) \leq \frac{30 w_i \log \frac{4k}{\delta}}{nw_i'}.$$

Since $w_i > \epsilon/4k$, with probability $\geq 1 - \delta/2k$, $w_i \leq 2w_i'$. By the union bound with probability $\geq 1 - \delta/k$, $w_i D(p_i, \hat{p}_i) \leq \frac{60 \log \frac{4k}{\delta}}{n}$. Hence if $n \geq \frac{120k \log \frac{4k}{\delta}}{\epsilon}$, the above quantity is less than $\epsilon/2k$. The total error probability is $\leq \delta$ by the union bound. □

Running MODIFIED SCHEFFE algorithm on the above set of candidates $\mathcal{F}$ yields a mixture that is close to the underlying one. By Lemma 1 and the above lemma we get

Corollary 11. *Let $n \geq c \cdot \frac{k \log \frac{k}{\epsilon\delta}}{\epsilon^2}$ for some constant c. There is an algorithm that runs in time*

$$\mathcal{O}\left(\left(\frac{k \log \frac{k}{\epsilon\delta}}{\epsilon}\right)^{3k-1} \frac{k^2 \log \frac{k}{\epsilon\delta}}{\epsilon^2}\right),$$

and returns a mixture $\hat{f}$ such that $D(f, \hat{f}) \leq 1000\epsilon$ with error probability $\leq 2\delta$.

Proof. Use $n' \stackrel{\text{def}}{=} \frac{120k \log \frac{4k}{\delta}}{\epsilon}$ samples to generate a set of at most n'^{3k-1} candidate distributions as stated in Lemma 10. With probability $\geq 1 - \delta$, one of the candidate distributions is ϵ-close to the underlying one. Run MODIFIED SCHEFFE on this set of candidate distributions to obtain a 1000ϵ-close estimate of f with probability $\geq 1 - \delta$ (Lemma 1). The run time is dominated by the run time of MODIFIED SCHEFFE which is $\mathcal{O}\left(\frac{|\mathcal{F}|T \log \frac{|\mathcal{F}|}{\delta}}{\epsilon^2}\right)$, where $|\mathcal{F}| = n'^{3k-1}$ and $T = k$. The total error probability is $\leq 2\delta$ by the union bound. □

Remark 12. *The above bound matches the independent and contemporary result by [15] for $k = 2$. While the process of identifying the candidate means is same for both the papers, the process of identifying the variances and proof techniques are different.*

5 Acknowledgements

We thank Sanjoy Dasgupta, Todd Kemp, and Krishnamurthy Vishwanathan for helpful discussions.

References

[1] Jayadev Acharya, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Optimal probability estimation with applications to prediction and classification. In *Proceedings of the 26th Annual Conference on Learning Theory (COLT)*, pages 764–796, 2013.

[2] Dimitris Achlioptas and Frank McSherry. On spectral learning of mixtures of distributions. In *Proceedings of the 18th Annual Conference on Learning Theory (COLT)*, pages 458–469, 2005.

[3] Rudolf Ahlswede and Andreas Winter. Strong converse for identification via quantum channels. *IEEE Transactions on Information Theory*, 48(3):569–579, 2002.

[4] Joseph Anderson, Mikhail Belkin, Navin Goyal, Luis Rademacher, and James R. Voss. The more, the merrier: the blessing of dimensionality for learning large gaussian mixtures. *CoRR*, abs/1311.2891, 2013.

[5] Martin Azizyan, Aarti Singh, and Larry A. Wasserman. Minimax theory for high-dimensional gaussian mixtures with sparse mean separation. *CoRR*, abs/1306.2035, 2013.

[6] Mikhail Belkin and Kaushik Sinha. Polynomial learning of distribution families. In *Proceedings of the 51st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 103–112, 2010.

[7] Kamalika Chaudhuri, Sanjoy Dasgupta, and Andrea Vattani. Learning mixtures of gaussians using the k-means algorithm. *CoRR*, abs/0912.0086, 2009.

[8] G.B. Coleman and Harry C. Andrews. Image segmentation by clustering. *Proceedings of the IEEE*, 67(5):773–785, 1979.

[9] Thomas M. Cover and Joy A. Thomas. *Elements of information theory (2. ed.)*. Wiley, 2006.

[10] Sanjoy Dasgupta. Learning mixtures of gaussians. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 634–644, 1999.

[11] Sanjoy Dasgupta and Leonard J. Schulman. A two-round variant of EM for gaussian mixtures. In *Proceedings of the 16th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 152–159, 2000.

[12] Sanjoy Dasgupta and Leonard J. Schulman. A probabilistic analysis of EM for mixtures of separated, spherical gaussians. *Journal on Machine Learning Research (JMLR)*, 8:203–226, 2007.

[13] Constantinos Daskalakis, Ilias Diakonikolas, and Rocco A. Servedio. Learning k-modal distributions via testing. In *SODA*, pages 1371–1385, 2012.

[14] Constantinos Daskalakis, Ilias Diakonikolas, and Rocco A. Servedio. Learning poisson binomial distributions. In *Proceedings of the 44th Annual Annual ACM Symposium on Theory of Computing (STOC)*, pages 709–728, 2012.

[15] Constantinos Daskalakis and Gautam Kamath. Faster and sample near-optimal algorithms for proper learning mixtures of gaussians. *CoRR*, abs/1312.1054, 2013.

[16] Luc Devroye and Gábor Lugosi. *Combinatorial methods in density estimation*. Springer, 2001.

[17] Inderjit S. Dhillon, Yuqiang Guan, and Jacob Kogan. Iterative clustering of high dimensional text data augmented by local search. In *Proceedings of the 2nd Industrial Conference on Data Mining (ICDM)*, pages 131–138, 2002.

[18] Jon Feldman, Ryan O'Donnell, and Rocco A. Servedio. Learning mixtures of product distributions over discrete domains. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 501–510, 2005.

[19] Jon Feldman, Rocco A. Servedio, and Ryan O'Donnell. PAC learning axis-aligned mixtures of gaussians with no separation assumption. In *Proceedings of the 19th Annual Conference on Learning Theory (COLT)*, pages 20–34, 2006.

[20] Yoav Freund and Yishay Mansour. Estimating a mixture of two product distributions. In *Proceedings of the 13th Annual Conference on Learning Theory (COLT)*, pages 53–62, 1999.

[21] Daniel Hsu and Sham M. Kakade. Learning mixtures of spherical gaussians: moment methods and spectral decompositions. In *Proceedings of the 4th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 11–20, 2013.

[22] Adam Tauman Kalai, Ankur Moitra, and Gregory Valiant. Efficiently learning mixtures of two gaussians. In *Proceedings of the 42nd Annual Annual ACM Symposium on Theory of Computing (STOC)*, pages 553–562, 2010.

[23] Ravindran Kannan, Hadi Salmasian, and Santosh Vempala. The spectral method for general mixture models. *SIAM Journal on Computing*, 38(3):1141–1156, 2008.

[24] Michael J. Kearns, Yishay Mansour, Dana Ron, Ronitt Rubinfeld, Robert E. Schapire, and Linda Sellie. On the learnability of discrete distributions. In *Proceedings of the 26th Annual Annual ACM Symposium on Theory of Computing (STOC)*, pages 273–282, 1994.

[25] B. Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *The Annals of Statistics*, 28(5):pp. 1302–1338, 2000.

[26] Bruce G. Lindsay. *Mixture Models: Theory, Geometry and Applications*. NSF-CBMS Conference series in Probability and Statistics, Penn. State University, 1995.

[27] Jinwen Ma, Lei Xu, and Michael I. Jordan. Asymptotic convergence rate of the em algorithm for gaussian mixtures. *Neural Computation*, 12(12):2881–2907, 2001.

[28] Satyaki Mahalanabis and Daniel Stefankovic. In *Proceedings of the 21st Annual Conference on Learning Theory (COLT)*, pages 503–512. Omnipress, 2008.

[29] Ankur Moitra and Gregory Valiant. Settling the polynomial learnability of mixtures of gaussians. In *Proceedings of the 51st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 93–102, 2010.

[30] Siu on Chan, Ilias Diakonikolas, Rocco A. Servedio, and Xiaorui Sun. Efficient density estimation via piecewise polynomial approximation. *CoRR*, abs/1305.3207, 2013.

[31] Siu on Chan, Ilias Diakonikolas, Rocco A. Servedio, and Xiaorui Sun. Learning mixtures of structured distributions over discrete domains. In *Proceedings of the 24th Annual Symposium on Discrete Algorithms (SODA)*, pages 1380–1394, 2013.

[32] Alon Orlitsky, Narayana P. Santhanam, Krishnamurthy Viswanathan, and Junan Zhang. On modeling profiles instead of values. In *Proceedings of the 20th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, 2004.

[33] Liam Paninski. Variational minimax estimation of discrete distributions under kl loss. In *Proceedings of the 18th Annual Conference on Neural Information Processing (NIPS)*, 2004.

[34] David Pollard. *Asymptopia*. 1997.

[35] Richard A. Redner and Homer F. Walker. Mixture densities, maximum likelihood and the em algorithm. *SIAM Review*, 26(2):pp. 195–239, 1984.

[36] Douglas A. Reynolds and Richard C. Rose. Robust text-independent speaker identification using gaussian mixture speaker models. *IEEE Transactions on Speech and Audio Processing*, 3(1):72–83, 1995.

[37] Robin Sibson. Slink: An optimally efficient algorithm for the single-link cluster method. *The Computer Journal*, 16(1):30–34, 1973.

[38] D Michael Titterington, Adrian FM Smith, and Udi E Makov. *Statistical analysis of finite mixture distributions*, volume 7. Wiley New York, 1985.

[39] Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.

[40] G. Valiant and P. Valiant. Estimating the unseen: an n/log(n)-sample estimator for entropy and support size, shown optimal via new clts. Proceedings of the 43rd Annual Annual ACM Symposium on Theory of Computing (STOC), 2011.

[41] Gregory Valiant and Paul Valiant. Estimating the unseen: A sublinear-sample canonical estimator of distributions. *Electronic Colloquium on Computational Complexity (ECCC)*, 17:180, 2010.

[42] Santosh Vempala and Grant Wang. A spectral algorithm for learning mixtures of distributions. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 113–, 2002.

[43] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *CoRR*, abs/1011.3027, 2010.

[44] Eric P. Xing, Michael I. Jordan, and Richard M. Karp. Feature selection for high-dimensional genomic microarray data. In *Proceedings of the 18th Annual International Conference on Machine Learning (ICML)*, pages 601–608, 2001.

[45] Bin Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer New York, 1997.

A Useful tools

A.1 Bounds on ℓ_1 distance

For two d dimensional product distributions $\mathbf{p}_1$ and $\mathbf{p}_2$, if we bound the ℓ_1 distance on each coordinate by ϵ, then by triangle inequality $D(\mathbf{p}_1, \mathbf{p}_2) \leq d\epsilon$. However this bound is often weak. One way to obtain a stronger bound is to relate ℓ_1 distance to Bhattacharyya parameter, which is defined as follows: Bhattacharyya parameter $B(p_1, p_2)$ between two distributions p_1 and p_2 is

$$B(p_1, p_2) = \int_{x \in \mathcal{X}} \sqrt{p_1(x) p_2(x)} dx.$$

We use the fact that for two product distributions $\mathbf{p}_1$ and $\mathbf{p}_2$, $B(\mathbf{p}_1, \mathbf{p}_2) = \prod_{i=1}^{d} B(p_{1,i}, p_{2,i})$ to obtain stronger bounds on the ℓ_1 distance. We first bound Bhattacharyya parameter for two one-dimensional Gaussian distributions.

Lemma 13. *The Bhattacharyya parameter for two one dimensional Gaussian distributions $p_1 = N(\mu_1, \sigma_1^2)$ and $p_2 = N(\mu_2, \sigma_2^2)$ is*

$$B(p_1, p_2) \geq 1 - \frac{(\mu_1 - \mu_2)^2)}{4(\sigma_1^2 + \sigma_2^2)} - \frac{(\sigma_1^2 - \sigma_2^2)^2}{(\sigma_1^2 + \sigma_2^2)^2}.$$

Proof. For Gaussian distributions the Bhattacharyya parameter is (see [8]), $B(p_1, p_2) = ye^{-x}$, where $x = \frac{(\mu_1 - \mu_2)^2)}{4(\sigma_1^2 + \sigma_2^2)}$ and $y = \sqrt{\frac{2\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2}}$. Observe that

$$y = \sqrt{\frac{2\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2}} = \sqrt{1 - \frac{(\sigma_1 - \sigma_2)^2}{\sigma_1^2 + \sigma_2^2}} \geq 1 - \frac{(\sigma_1 - \sigma_2)^2}{\sigma_1^2 + \sigma_2^2} \geq 1 - \frac{(\sigma_1^2 - \sigma_2^2)^2}{(\sigma_1^2 + \sigma_2^2)^2}.$$

Hence,

$$B(p_1, p_2) = ye^{-x} \geq y(1 - x) \geq (1 - x)\left(1 - \frac{(\sigma_1^2 - \sigma_2^2)^2}{(\sigma_1^2 + \sigma_2^2)^2}\right) \geq 1 - x - \frac{(\sigma_1^2 - \sigma_2^2)^2}{(\sigma_1^2 + \sigma_2^2)^2}.$$

Substituting the value of x results in the lemma. □

The next lemma follows from the relationship between Bhattacharyya parameter and ℓ_1 distance (see [34]), and the previous lemma.

Lemma 14. *For any two Gaussian product distributions $\mathbf{p}_1$ and $\mathbf{p}_2$,*

$$D(\mathbf{p}_1, \mathbf{p}_2)^2 \leq 8\left(\sum_{i=1}^{d} 1 - B(p_{1,i}, p_{2,i})\right) \leq \sum_{i=1}^{d} 2\frac{(\mu_{1,i} - \mu_{2,i})^2}{\sigma_{1,i}^2 + \sigma_{2,i}^2} + 8\frac{(\sigma_{1,i}^2 - \sigma_{2,i}^2)^2}{(\sigma_{1,i}^2 + \sigma_{2,i}^2)^2}.$$

A.2 Concentration inequalities

We use the following concentration inequalities for Gaussian, Chi-Square, and sum of Bernoulli random variables in the rest of the paper.

Lemma 15. *For a Gaussian random variable X with mean μ and variance σ^2,*

$$\Pr(|X - \mu| \geq t\sigma) \leq e^{-t^2/2}.$$

Lemma 16 ([25])**.** *If $Y_1, Y_2, \ldots Y_n$ be n i.i.d.Gaussian variables with mean 0 and variance σ^2, then*

$$\Pr\left(\sum_{i=1}^{n} Y_i^2 - n\sigma^2 \geq 2(\sqrt{nt} + t)\sigma^2\right) \leq e^{-t}, \text{ and } \Pr\left(\sum_{i=1}^{n} Y_i^2 - n\sigma^2 \leq -2\sqrt{nt}\sigma^2\right) \leq e^{-t}.$$

Furthermore for a fixed vector $\mathbf{a}$,

$$\Pr\left(\left|\sum_{i=1}^{n} \mathbf{a}_i(Y_i^2 - 1)\right| \leq 2(\|\mathbf{a}\|_2 \sqrt{t} + \|\mathbf{a}\|_\infty t)\sigma^2\right) \leq 2e^{-t}.$$

Lemma 17 (Chernoff bound)**.** *If $X_1, X_2 \ldots X_n$ are distributed according to Bernoulli p, then with probability $1 - \delta$,*

$$\left|\frac{\sum_{i=1}^{n} X_i}{n} - p\right| \leq \sqrt{\frac{2p(1-p)}{n}\log\frac{2}{\delta}} + \frac{2}{3}\frac{\log\frac{2}{\delta}}{n}.$$

We now state a non-asymptotic concentration inequality for random matrices that helps us bound errors in spectral algorithms.

Lemma 18 ([43] Remark 5.51)**.** *Let $\mathbf{y}(1), \mathbf{y}(2), \ldots, \mathbf{y}(n)$ be generated according to $N(0, \Sigma)$. For every $\epsilon \in (0, 1)$ and $t \geq 1$, if $n \geq c' d\left(\frac{t}{\epsilon}\right)^2$ for some constant c', then with probability $\geq 1 - 2e^{-t^2 n}$,*

$$\left\|\sum_{i=1}^{n} \frac{1}{n}\mathbf{y}(i)\mathbf{y}^t(i) - \Sigma\right\| \leq \epsilon \|\Sigma\|.$$

A.3 Matrix eigenvalues

We now state few simple lemmas on the eigenvalues of perturbed matrices.

Lemma 19. *Let $\lambda_1^A \geq \lambda^A \geq \ldots \lambda_d^A \geq 0$ and $\lambda_1^B \geq \lambda^B \geq \ldots \lambda_d^B \geq 0$ be the eigenvalues of two symmetric matrices A and B respectively. If $\|A - B\| \leq \epsilon$, then $\forall\, i$, $|\lambda_i^A - \lambda_i^B| \leq \epsilon$.*

Proof. Let $\mathbf{u}_1, \mathbf{u}_2, \ldots \mathbf{u}_d$ be a set of eigenvectors of A that corresponds to $\lambda_1^A, \lambda_2^A, \ldots \lambda_d^A$. Similarly let $\mathbf{v}_1, \mathbf{v}_2, \ldots \mathbf{v}_d$ be eigenvectors of B Consider the first eigenvalue of B,

$$\lambda_1^B = \|B\| = \|A + (B - A)\| \geq \|A\| - \|B - A\| \geq \lambda_1^A - \epsilon.$$

Now consider an $i > 1$. If $\lambda_i^B < \lambda_i^A - \epsilon$, then by definition of eigenvalues

$$\max_{\mathbf{v}:\forall j \le i-1, \mathbf{v}\cdot\mathbf{v}_j=0} \|B\mathbf{v}\|_2 < \lambda_i^A - \epsilon.$$

Now consider a unit vector $\sum_{j=1}^{i} \alpha_j \mathbf{u}_j$ in the span of $\mathbf{u}_1, \ldots \mathbf{u}_i$, that is orthogonal to $\mathbf{v}_1, \ldots \mathbf{v}_{i-1}$. For this vector,

$$\left\| B \sum_{j=1}^{i} \alpha_j \mathbf{u}_j \right\|_2 \ge \left\| A \sum_{j=1}^{i} \alpha_j \mathbf{u}_j \right\|_2 - \left\| (A-B) \sum_{j=1}^{i} \alpha_j \mathbf{u}_j \right\|_2 \ge \sqrt{\sum_{j=1}^{i} \alpha_j^2 (\lambda_j^A)^2} - \epsilon \ge \lambda_i^A - \epsilon,$$

a contradiction. Hence, $\forall i \le d$, $\lambda_i^B \ge \lambda_i^A - \epsilon$. The proof in the other direction is similar and omitted. □

Lemma 20. *Let $A = \sum_{i=1}^{k} \eta_i^2 \mathbf{u}_i \mathbf{u}_i^t$ be a positive semidefinite symmetric matrix for $k \le d$. Let $\mathbf{u}_1, \mathbf{u}_2, \ldots \mathbf{u}_k$ span a $k-1$ dimensional space. Let $B = A + R$, where $\|R\| \le \epsilon$. Let $\mathbf{v}_1, \mathbf{v}_2, \ldots \mathbf{v}_{k-1}$ be the top $k-1$ eigenvectors of B. Then the projection of $\mathbf{u}_i$ in space orthogonal to $\mathbf{v}_1, \mathbf{v}_2, \ldots \mathbf{v}_{k-1}$ is $\le \frac{2\sqrt{\epsilon}}{\eta_i}$.*

Proof. Let λ_i^B be the i^{th} largest eigenvalue of B. Observe that $B + \epsilon \mathbb{I}_d$ is a positive semidefinite matrix as for any vector $\mathbf{v}$, $\mathbf{v}^t(A + R + \epsilon \mathbb{I}_d)\mathbf{v} \ge 0$. Furthermore $\|A + R + \epsilon \mathbb{I}_d - A\| \le 2\epsilon$. Since eigenvalues of $B + \epsilon \mathbb{I}_d$ is $\lambda^B + \epsilon$, by Lemma 19, for all $i \le d$, $|\lambda_i^A - \lambda_i^B - \epsilon| \le 2\epsilon$. Therefore, $|\lambda_i^B|$ for $i \ge k$ is $\le 3\epsilon$.

Let $\mathbf{u}_i = \sum_{j=1}^{k-1} \alpha_{i,j} \mathbf{v}_j + \sqrt{1 - \sum_{j=1}^{k-1} \alpha_{i,j}^2}\, \mathbf{u}'$, for a vector $\mathbf{u}'$ orthogonal to $\mathbf{v}_1, \mathbf{v}_2, \ldots \mathbf{v}_{k-1}$. We compute $\mathbf{u}'^t A \mathbf{u}'$ in two ways. Since $A = B - R$,

$$|\mathbf{u}'^t (B-R)\mathbf{u}'| \le |\mathbf{u}'^t B \mathbf{u}'| + |\mathbf{u}'^t R \mathbf{u}'| \le \left\|B\mathbf{u}'\right\|_2 + \|R\|.$$

Since $\mathbf{u}'$ is orthogonal to first k eigenvectors, we have $\|B\mathbf{u}'\|_2 \le 3\epsilon$ and hence $|\mathbf{u}'^(B - R)\mathbf{u}'| \le 4\epsilon$.

$$\mathbf{u}'^t A \mathbf{u}' \ge \eta_i^2 \big(1 - \sum_{j=1}^{k-1} \alpha_{i,j}^2\big).$$

We have shown that the above quantity is $\le 4\epsilon$. Therefore $\left(1 - \sum_{j=1}^{k-1} \alpha_{i,j}^2\right)^{1/2} \le 2\sqrt{\epsilon}/\eta_i$. □

B Selection from a set of candidate distributions

Given samples from an unknown distribution f, the objective is to output a distribution from a known collection $\mathcal{F}$ of distributions with ℓ_1 distance close to $D(f, \mathcal{F})$. Scheffe estimate [16] outputs a distribution from $\mathcal{F}$ whose ℓ_1 distance from f is at most $9.1 \max(D(f, \mathcal{F}), \epsilon)$ The algorithm requires $\mathcal{O}(\epsilon^{-2} \log |\mathcal{F}|)$ samples and the runs in time $\mathcal{O}(|\mathcal{F}|^2 T(n + |\mathcal{X}|))$, where T is the time to compute the probability $f_j(x)$ of x, for any $f_j \in \mathcal{F}$. An approach to reduce the time complexity, albeit using exponential pre-processing, was proposed in [28]. We present the modified Scheffe algorithm with near linear time complexity and then prove Lemma 1.

We first present the algorithm SCHEFFE* with running time $\widetilde{\mathcal{O}}(|\mathcal{F}|^2 T n)$.

Algorithm SCHEFFE*
Input: a set $\mathcal{F}$ of candidate distributions, ϵ : upper bound on $D(f, \mathcal{F})$, n independent samples $x_1, \ldots, x_n$ from f.

For each pair (p, q) in $\mathcal{F}$ do:

1. $\mu_f = \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}\{p(x_i) > q.(x_i)\}$.
2. Generate independent samples $y_1, \ldots, y_n$ and $z_1, \ldots, z_n$ from p and q respectively.
3. $\mu_p = \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}\{p(y_i) > q(y_i)\}$, $\mu_q = \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}\{p(z_i) > q(z_i)\}$.
4. If $|\mu_p - \mu_f| < |\mu_q - \mu_f|$ declare p as winner, else q.

Output the distribution with most wins, breaking ties arbitrarily.

We make the following modification to the algorithm where we reduce the size of potential distributions by half in every iteration.

Algorithm MODIFIED SCHEFFE
Input: set $\mathcal{F}$ of candidate distributions, ϵ : upper bound on $\min_{f_i \in \mathcal{F}} D(f, f_i)$, n independent samples $x_1, \ldots, x_n$ from f.

1. Let $\mathcal{G} = \mathcal{F}$, $\mathcal{C} \leftarrow \varnothing$
2. Repeat until $|\mathcal{G}| > 1$:
   (a) Randomly form $|\mathcal{G}|/2$ pairs of distributions in $\mathcal{G}$ and run SCHEFFE* on *each pair* using the n samples.
   (b) Replace $\mathcal{G}$ with the $|\mathcal{G}|/2$ winners.
   (c) Randomly select a set $\mathcal{A}$ of $\min\{|\mathcal{G}|, |\mathcal{F}|^{1/3}\}$ elements from $\mathcal{G}$.
   (d) Run SCHEFFE* on each pair in $\mathcal{A}$ and add the distributions with most wins to $\mathcal{C}$.
3. Run SCHEFFE* on $\mathcal{C}$ and output the winner

Remark 21. *For the ease of proof, we assume that $\delta \geq \frac{10 \log |\mathcal{F}|}{|\mathcal{F}|^{1/3}}$. If $\delta < \frac{10 \log |\mathcal{F}|}{|\mathcal{F}|^{1/3}}$, we run the algorithm with error probability $1/3$ and repeat it $\mathcal{O}(\log \frac{1}{\delta})$ times to choose a set of candidate mixtures $\mathcal{F}_\delta$. By Chernoff-bound with error probability $\leq \delta$, $\mathcal{F}_\delta$ contains a mixture close to f. Finally, we run* SCHEFFE* *on $\mathcal{F}_\delta$ to obtain a mixture that is close to f.*

Proof sketch of Lemma 1. For any set $\mathcal{A}$ and a distribution p, given n independent samples from p the empirical probability $\mu_n(\mathcal{A})$ has a distribution around $p(\mathcal{A})$ with standard deviation $\sim \frac{1}{\sqrt{n}}$. Together with an observation in Scheffe estimation in [16] one can show that if the number of samples $n = \mathcal{O}\left(\frac{\log \frac{|\mathcal{F}|}{\delta}}{\epsilon^2}\right)$, then SCHEFFE* has a guarantee $10 \max(\epsilon, D(f, \mathcal{F}))$ with probability $\geq 1 - \delta$.

Since we run SCHEFFE* at most $|\mathcal{F}|(2\log|\mathcal{F}|+1)$ times, choosing $\delta = \delta/(4|\mathcal{F}|\log|\mathcal{F}|+2|\mathcal{F}|)$ results in the sample complexity of

$$\mathcal{O}\left(\frac{\log\frac{|\mathcal{F}|^2(4\log|\mathcal{F}|+2)}{\delta}}{\epsilon^2}\right) = \mathcal{O}\left(\frac{\log\frac{|\mathcal{F}|}{\delta}}{\epsilon^2}\right),$$

and the total error probability of $\delta/2$ for all runs of SCHEFFE* during the algorithm. The above value of n dictates our sample complexity. We now consider the following two cases:

- If at some stage $\geq \frac{\log(2/\delta)}{|\mathcal{F}|^{1/3}}$ fraction of elements in $\mathcal{A}$ have an ℓ_1 distance $\leq 10\epsilon$ from f, then at that stage with probability $\geq 1-\delta/2$ an element with distance $\leq 10\epsilon$ from f is added to $\mathcal{A}$. Therefore a distribution with distance $\leq 100\epsilon$ is selected to $\mathcal{C}$.
- If at no stage this happens, then consider the element that is closest to f, *i.e.*, at ℓ_1 distance at most ϵ. With probability $\geq \left(1-\frac{\log(2/\delta)}{|\mathcal{F}|^{1/3}}\right)^{\log|\mathcal{F}|}$ it always competes with an element at a distance at least 10ϵ from f and it wins all these games with probability $\geq 1-\delta/2$.

Therefore with probability $\geq 1-\delta/2$ there is an element in $\mathcal{C}$ at ℓ_1 distance at most 100ϵ. Running SCHEFFE* on this set yields a distribution at a distance $\leq 100\cdot 10\epsilon = 1000\epsilon$. The error probability is $\leq \delta$ by the union bound. □

C Lower bound

We first show a lower bound for a single Gaussian distribution and generalize it to mixtures.

C.1 Single Gaussian distribution

The proof is an application of the following version of Fano's inequality [9, 45]. It states that we cannot simultaneously estimate all distributions in a class using n samples if they satisfy certain conditions.

Lemma 22. *(Fano's Inequality) Let $f_1,\ldots,f_{r+1}$ be a collection of distributions such that for any $i\neq j$, $D(f_i,f_j)\geq\alpha$, and $KL(f_i,f_j)\leq\beta$. Let f be an estimate of the underlying distribution using n i.i.d. samples from one of the f_i's. Then,*

$$\sup_i \mathbb{E}[D(f_i,f)] \geq \frac{\alpha}{2}\left(1-\frac{n\beta+\log 2}{\log r}\right).$$

We consider $d-$dimensional spherical Gaussians with identity covariance matrix, with means along any coordinate restricted to $\pm\frac{c\epsilon}{\sqrt{d}}$. The KL divergence between two spherical Gaussians with identity covariance matrix is the squared distance between their means. Therefore, any two distributions we consider have KL distance at most

$$\beta = \sum_{i=1}^{d}\left(2\frac{c\epsilon}{\sqrt{d}}\right)^2 = 4c^2\epsilon^2,$$

We now consider a subset of these 2^d distributions to obtain a lower bound on α. By the Gilbert-Varshamov bound, there exists a binary code with $\geq 2^{d/8}$ codewords of length d and minimum distance $d/8$. Consider one such code. Now for each codeword, map $1\to\frac{c\epsilon}{\sqrt{d}}$ and $0\to-\frac{c\epsilon}{\sqrt{d}}$ to obtain a distribution in our class. We consider this subset of $\geq 2^{d/8}$ distributions as our f_i's.

Consider any two f_i's. Their means differ in at least $d/8$ coordinates. We show that the ℓ_1 distance between them is $\geq c\epsilon/4$. Without loss of generality, let the means differ in the first $d/8$ coordinates, and furthermore, one of the distributions has means $c\epsilon/\sqrt{d}$ and the other has $-c\epsilon/\sqrt{d}$ in the first $d/8$ coordinates. The sum of the first $d/8$ coordinates is $N(c\epsilon\sqrt{d}/8, d/8)$ and $N(-c\epsilon\sqrt{d}/8, d/8)$. The ℓ_1 distance between these normal random variables is a lower bound on the ℓ_1 distance of the original random variables. For small values of $c\epsilon$ the distance between the two Gaussians is at least $\geq c\epsilon/4$. This serves as our α.

Applying the Fano's Inequality, the ℓ_1 error on the worst distribution is at least

$$\frac{c\epsilon}{8}\Big(1 - \frac{n4c^2\epsilon^2 + \log 2}{d/8}\Big),$$

which for $c = 16$ and $n < \frac{d}{2^{14}\epsilon^2}$ is at least ϵ. In other words, the smallest n to approximate all spherical normal distributions to ℓ_1 distance at most ϵ is $> \frac{d}{2^{14}\epsilon^2}$.

C.2 Mixtures of k Gaussians

We now provide a lower bound on the sample complexity of learning mixtures of k Gaussians in d dimensions. We extend the construction for learning a single spherical Gaussian to mixtures of k Gaussians and show a lower bound of $\Omega(kd/\epsilon^2)$ samples. We will again use Fano's inequality over a class of $2^{kd/64}$ distributions as described next.

To prove the lower bound on the sample complexity of learning spherical Gaussians, we designed a class of $2^{d/8}$ distributions around the origin. Let $\mathcal{P} \stackrel{\text{def}}{=} \{P_1, \ldots, P_T\}$, where $T = 2^{d/8}$, be this class. Recall that each P_i is a spherical Gaussian with unit variance. For a distribution P over $\mathbb{R}^d$ and $\boldsymbol{\mu} \in \mathbb{R}^d$, let $P + \boldsymbol{\mu}$ be the distribution P shifted by $\boldsymbol{\mu}$.

We now choose $\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_k$'s *extremely well-separated*. The class of distributions we consider will be a mixture of k components, where the jth component is a distribution from $\mathcal{P}$ shifted by $\boldsymbol{\mu}_j$. Since the $\boldsymbol{\mu}$'s will be well separated, we will use the results from last section over each component.

For $i \in [T]$, and $j \in [k]$, $P_{ij} \stackrel{\text{def}}{=} P_i + \boldsymbol{\mu}_j$. Each $(i_1, \ldots, i_k) \in [T]^k$ corresponds to the mixture

$$\frac{1}{k}(P_{i_1 1} + P_{i_2 2} + \ldots + P_{i_k k})$$

of k spherical Gaussians. We consider this class of $T^k = 2^{kd/8}$ distributions. By the Gilbert-Varshamov bound, for any $T \geq 2$, there is a T-ary codes of length k, with minimum distance $\geq k/8$ and number of codewords $\geq 2^{k/8}$. This implies that among the $T^k = 2^{dk/8}$ distributions, there are $2^{kd/64}$ distributions such that any two tuples $(i_1, \ldots, i_k)$ and $(i'_1, \ldots, i'_k)$ corresponding to different distributions differ in at least $k/8$ locations.

If we choose the $\boldsymbol{\mu}$'s well separated, the components of any mixture distribution have very little overlap. For simplicity, we choose $\boldsymbol{\mu}_j$'s satisfying

$$\min_{j_1 \neq j_2} \|\boldsymbol{\mu}_{j_1} - \boldsymbol{\mu}_{j_2}\|_2 \geq \left(\frac{2kd}{\epsilon}\right)^{100}.$$

This implies that for $j \neq l$, $\|P_{ij} - P_{i'l}\|_1 < (\epsilon/2dk)^{10}$. Therefore, for two different mixture distributions,

$$\begin{aligned}
&\left\| \frac{1}{k}(P_{i_1 1} + P_{i_2 2} + \ldots + P_{i_k k}) - \frac{1}{k}(P_{i'_1 1} + P_{i'_2 2} + \ldots + P_{i'_k k}) \right\|_1 \\
&\overset{(a)}{\geq} \frac{1}{k} \sum_{j \in [k], i_j, i'_j \in [T]} |P_{i_j j} - P_{i'_j j}| - k^2 (\epsilon/2dk)^{10} \\
&\overset{(b)}{\geq} \frac{1}{8}\frac{c\epsilon}{4} - k^2(\epsilon/2dk)^{10}.
\end{aligned}$$

where (a) follows form the fact that two mixtures have overlap only in the corresponding components, (b) uses the fact that at least in $k/8$ components $i_j \neq i'_j$, and then uses the lower bound from the previous section.

Therefore, the ℓ_1 distance between any two of the $2^{kd/64}$ distributions is $\geq c_1\epsilon/32$ for c_1 slightly smaller than c. We take this as α.

Now, to upper bound the KL divergence, we simply use the convexity, namely for any distributions $P_1 \ldots P_k$ and $Q_1 \ldots Q_k$, let $\bar{P}$ and $\bar{Q}$ be the mean distributions. Then,

$$D(\bar{P}\|\bar{Q}) \leq \frac{1}{k}\sum_{i=1}^{k} D(P_i\|Q_i).$$

By the construction and from the previous section, for any j,

$$D(P_{i_j j}\|P_{i'_j j}) = D(P_i\|P_{i'}) \leq 4c^2\epsilon^2.$$

Therefore, we can take $\beta = 4c^2\epsilon^2$.

Therefore by the Fano's inequality, the ℓ_1 error on the worst distribution is at least

$$\frac{c_1\epsilon}{64}\Big(1 - \frac{n4c^2\epsilon^2 + \log 2}{dk/64}\Big),$$

which for $c_1 = 128, c = 128.1$ and $n < \frac{dk}{8^8\epsilon^2}$ is at least ϵ.

D Proofs for k spherical Gaussians

We first state a simple concentration result that helps us in other proofs.

Lemma 23. *Given n samples from a set of Gaussian distributions, with probability $\geq 1 - 2\delta$, for every pair of samples $\mathbf{X} \sim N(\boldsymbol{\mu}_1, \sigma^2\mathbb{I}_d)$ and $\mathbf{Y} \sim N(\boldsymbol{\mu}_2, \sigma^2\mathbb{I}_d)$,*

$$\|\mathbf{X} - \mathbf{Y}\|_2^2 \leq 2d\sigma^2 + 4\sigma^2\sqrt{d\log\frac{n^2}{\delta}} + \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2 + 4\sigma\|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2\sqrt{\log\frac{n^2}{\delta}} + 4\sigma^2\log\frac{n^2}{\delta}. \quad (1)$$

and

$$\|\mathbf{X} - \mathbf{Y}\|_2^2 \geq 2d\sigma^2 - 4\sigma^2\sqrt{d\log\frac{n^2}{\delta}} + \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2 - 4\sigma\|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2\sqrt{\log\frac{n^2}{\delta}}. \quad (2)$$

Proof. We prove the lower bound, the proof for the upper bound is similar and omitted. Since $\mathbf{X}$ and $\mathbf{Y}$ are Gaussians, $\mathbf{X} - \mathbf{Y}$ is distributed as $N(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2, 2\sigma^2)$. Rewriting $\|\mathbf{X} - \mathbf{Y}\|_2$

$$\|\mathbf{X} - \mathbf{Y}\|_2^2 = \|\mathbf{X} - \mathbf{Y} - (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)\|_2^2 + \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2 + 2(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \cdot (\mathbf{X} - \mathbf{Y} - (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)).$$

Let $\mathbf{Z} = \mathbf{X} - \mathbf{Y} - (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)$, then $\mathbf{Z} \sim N(\mathbf{0}, 2\sigma^2 \mathbb{I}_d)$. Therefore by Lemma 16, with probability $1 - \delta/n^2$,

$$\|\mathbf{Z}\|_2^2 \geq 2d\sigma^2 - 4\sigma^2 \sqrt{d \log \frac{n^2}{\delta}}.$$

Furthermore $(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \cdot \mathbf{Z}$ is sum of Gaussians and hence a Gaussian distribution. It has mean 0 and variance $2\sigma^2 \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2$. Therefore, by Lemma 15 with probability $1 - \delta/n^2$,

$$(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \cdot \mathbf{Z} \geq -2\sigma \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2 \sqrt{\log \frac{n^2}{\delta}}.$$

By the union bound with probability $1 - 2\delta/n^2$,

$$\|\mathbf{X} - \mathbf{Y}\|_2^2 \geq 2d\sigma^2 - 4\sigma^2 \sqrt{d \log \frac{n^2}{\delta}} + \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2 - 4\sigma \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2 \sqrt{\log \frac{n^2}{\delta}}.$$

There are $\binom{n}{2}$ pairs and the lemma follows by the union bound. □

D.1 Proof of Lemma 4

We show that if Equations (1) and (2) are satisfied, then the lemma holds. The error probability is that of Lemma 23 and is $\leq 2\delta$. Since the minimum is over $k + 1$ indices, at least two samples are from the same component. Applying Equations (1) and (2) for these two samples

$$2d\hat{\sigma}^2 \leq 2d\sigma^2 + 4\sigma^2 \sqrt{d \log \frac{n^2}{\delta}} + 4\sigma^2 \log \frac{n^2}{\delta}.$$

Similarly by Equations (1) and (2) for any two samples $\mathbf{X}(a), \mathbf{X}(b)$ in $[k + 1]$,

$$\begin{aligned}\|\mathbf{X}(a) - \mathbf{X}(b)\|_2^2 &\geq 2d\sigma^2 - 4\sigma^2 \sqrt{d \log \frac{n^2}{\delta}} + \left\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\right\|_2^2 - 4\sigma \left\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\right\|_2 \sqrt{\log \frac{n^2}{\delta}} \\ &\geq 2d\sigma^2 - 4\sigma^2 \sqrt{d \log \frac{n^2}{\delta}} - 4\sigma^2 \log \frac{n^2}{\delta},\end{aligned}$$

where the last inequality follows from the fact that $\alpha^2 - 4\alpha\beta \geq -4\beta^2$. The result follows from the assumption that $d > 20 \log n^2/\delta$.

D.2 Proof of Lemma 5

We show that if Equations (1) and (2) are satisfied, then the lemma holds. The error probability is that of Lemma 23 and is $\leq 2\delta$. Since Equations (1) and (2) are satisfied, by the proof of Lemma 4, $|\hat{\sigma}^2 - \sigma^2| \leq 2.5\sigma^2 \sqrt{\frac{\log(n^2/\delta)}{d}}$. If two samples $X(a)$ and $X(b)$ are from the same component, by Lemma 23,

$$\|\mathbf{X}(a) - \mathbf{X}(b)\|_2^2 \leq 2d\sigma^2 + 4\sigma^2 \sqrt{d \log \frac{n^2}{\delta}} + 4\sigma^2 \log \frac{n^2}{\delta} \leq 2d\sigma^2 + 5\sigma^2 \sqrt{d \log \frac{n^2}{\delta}}.$$

By Lemma 4, the above quantity is less than $2d\hat{\sigma}^2 + 23\hat{\sigma}^2\sqrt{d\log\frac{n^2}{\delta}}$. Hence all the samples from the same component are in a single cluster.

Suppose there are two samples from different components in a cluster, then by Equations (1) and (2),

$$2d\hat{\sigma}^2 + 23\hat{\sigma}^2\sqrt{d\log\frac{n^2}{\delta}} \geq 2d\sigma^2 - 4\sigma^2\sqrt{d\log\frac{n^2}{\delta}} + \left\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\right\|_2^2 - 4\sigma\left\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\right\|_2\sqrt{\log\frac{n^2}{\delta}}.$$

Relating $\hat{\sigma}^2$ and σ^2 using Lemma 4,

$$2d\sigma^2 + 40\sigma^2\sqrt{d\log\frac{n^2}{\delta}} \geq 2d\sigma^2 - 4\sigma^2\sqrt{d\log\frac{n^2}{\delta}} + \left\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\right\|_2^2 - 4\sigma\left\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\right\|_2\sqrt{\log\frac{n^2}{\delta}}.$$

Hence $\left\|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\right\|_2 \leq 10\sigma\left(d\log\frac{n^2}{\delta}\right)^{1/4}$. There are at most k components; therefore, any two components within the same cluster are at a distance $\leq 10k\sigma\left(d\log\frac{n^2}{\delta}\right)^{1/4}$.

D.3 Proof of Lemma 6

The proof is involved and we show it in steps. We first show few concentration bounds which we use later to argue that the samples are clusterable when the sample covariance matrix has a large eigenvalue. Let $\hat{w}_i$ be the fraction of samples from component i. Let $\hat{\boldsymbol{\mu}}_i$ be the empirical average of samples from $\mathbf{p}_i$. Let $\hat{\overline{\boldsymbol{\mu}}}(C)$ be the empirical average of samples in cluster C. If C is the entire set of samples we use $\hat{\overline{\boldsymbol{\mu}}}$ instead of $\hat{\overline{\boldsymbol{\mu}}}(C)$. We first show a concentration inequality that we use in rest of the calculations.

Lemma 24. *Given n samples from a k-component Gaussian mixture with probability $\geq 1 - 2\delta$, for every component i*

$$\|\hat{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i\|_2^2 \leq \left(d + 3\sqrt{d\log\frac{2k}{\delta}}\right)\frac{\sigma^2}{n\hat{w}_i} \text{ and } |\hat{w}_i - w_i| \leq \sqrt{\frac{2w_i\log\frac{2k}{\delta}}{n}} + \frac{2}{3}\frac{\log\frac{2k}{\delta}}{n}. \tag{3}$$

Proof. Since $\hat{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i$ is distributed $N(0, \sigma^2\mathbb{I}_d/n\hat{w}_i)$, by Lemma 16 with probability $\geq 1 - \delta/k$,

$$\|\hat{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i\|_2^2 \leq \left(d + 2\sqrt{d\log\frac{2k}{\delta}} + 2\log\frac{2k}{\delta}\right)\frac{\sigma^2}{n\hat{w}_i} \leq \left(d + 3\sqrt{d\log\frac{2k}{\delta}}\right)\frac{\sigma^2}{n\hat{w}_i}.$$

The second inequality uses the fact that $d \geq 20\log n^2/\delta$. For bounding the weights, observe that by Lemma 17 with probability $\geq 1 - \delta/k$,

$$|\hat{w}_i - w_i| \leq \sqrt{\frac{2w_i\log 2k/\delta}{n}} + \frac{2}{3}\frac{\log 2k/\delta}{n}.$$

By the union bound the error probability is $\leq 2k\delta/2k = \delta$. □

A simple application of triangle inequality yields the following lemma.

Lemma 25. *Given n samples from a k-component Gaussian mixture if Equation* (3) *holds, then*

$$\left\|\sum_{i=1}^{k}\hat{w}_i(\hat{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i)(\hat{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i)^t\right\| \leq \left(d + 3\sqrt{d\log\frac{2k}{\delta}}\right)\frac{k\sigma^2}{n}.$$

Lemma 26. *Given n samples from a k-component Gaussian mixture, if Equation* (3) *holds and the maximum distance between two components is $\le 10k\sigma\left(d\log\frac{n^2}{\delta}\right)^{1/4}$, then $\left\|\hat{\overline{\boldsymbol{\mu}}}-\overline{\boldsymbol{\mu}}\right\|_2 \le c\sigma\sqrt{\frac{dk\log\frac{n^2}{\delta}}{n}}$, for a constant c.*

Proof. Observe that

$$\hat{\overline{\boldsymbol{\mu}}}-\overline{\boldsymbol{\mu}} = \sum_{i=1}^{k}\hat{w}_i\hat{\boldsymbol{\mu}}_i - w_i\boldsymbol{\mu}_i = \sum_{i=1}^{k}\hat{w}_i(\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i)+(\hat{w}_i-w_i)\boldsymbol{\mu}_i = \sum_{i=1}^{k}\hat{w}_i(\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i)+(\hat{w}_i-w_i)(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}). \quad (4)$$

Hence by Equation (3) and the fact that the maximum distance between two components is $\le 10k\sigma\left(d\log\frac{n^2}{\delta}\right)^{1/4}$,

$$\left\|\hat{\overline{\boldsymbol{\mu}}}-\overline{\boldsymbol{\mu}}\right\|_2 \le \sum_{i=1}^{k}\hat{w}_i\sqrt{\left(d+3\sqrt{d\log\frac{2k}{\delta}}\right)}\frac{\sigma}{\sqrt{n\hat{w}_i}} + \left(\sqrt{\frac{2w_i\log 2k/\delta}{n}}+\frac{2}{3}\frac{\log 2k/\delta}{n}\right)10k\left(d\log\frac{n^2}{\delta}\right)^{1/4}\sigma.$$

For $n \ge d \ge \max(k^4, 20\log n^2/\delta, 1000)$, we get the above term is $\le c\sqrt{\frac{kd\log n^2/\delta}{n}}\sigma$, for some constant c. □

We now make a simple observation on covariance matrices.

Lemma 27. *Given n samples from a k-component mixture,*

$$\begin{aligned}&\left\|\sum_{i=1}^{k}\hat{w}_i(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})^t - \sum_{i=1}^{k}\hat{w}_i(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})^t\right\| \\ &\le 2\left\|\hat{\overline{\boldsymbol{\mu}}}-\overline{\boldsymbol{\mu}}\right\|_2^2 + \sum_{i=1}^{k}2\hat{w}_i\left\|\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i\right\|_2^2 + 2\left(\sqrt{k}\left\|\hat{\overline{\boldsymbol{\mu}}}-\overline{\boldsymbol{\mu}}\right\|_2 + \sum_{i=1}^{k}\sqrt{\hat{w}_i}\left\|\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i\right\|_2\right)\max_j\sqrt{\hat{w}_j}\left\|\boldsymbol{\mu}_j-\overline{\boldsymbol{\mu}}\right\|_2.\end{aligned}$$

Proof. Observe that for any two vectors $\mathbf{u}$ and $\mathbf{v}$,

$$\mathbf{u}\mathbf{u}^t-\mathbf{v}\mathbf{v}^t = \mathbf{u}(\mathbf{u}^t-\mathbf{v}^t)+(\mathbf{u}-\mathbf{v})\mathbf{v}^t = (\mathbf{u}-\mathbf{v})(\mathbf{u}-\mathbf{v})^t+\mathbf{v}(\mathbf{u}-\mathbf{v})^t+(\mathbf{u}-\mathbf{v})\mathbf{v}^t.$$

Hence by triangle inequality,

$$\left\|\mathbf{u}\mathbf{u}^t-\mathbf{v}\mathbf{v}^t\right\| \le \|\mathbf{u}-\mathbf{v}\|_2^2 + 2\|\mathbf{v}\|_2\|\mathbf{u}-\mathbf{v}\|_2.$$

Applying the above observation to $\mathbf{u} = \hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}}$ and $\mathbf{v} = \boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}$, we get

$$\begin{aligned}&\sum_{i=1}^{k}\hat{w}_i\left\|(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})^t-(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})^t\right\| \\ &\le \sum_{i=1}^{k}\left(\hat{w}_i\left\|\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}}-\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}\right\|_2^2 + 2\sqrt{\hat{w}_i}\left\|\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}\right\|_2\sqrt{\hat{w}_i}\left\|\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}}-\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}\right\|_2\right) \\ &\le \sum_{i=1}^{k}\left(2\hat{w}_i\left\|\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i\right\|_2^2 + 2\hat{w}_i\left\|\hat{\overline{\boldsymbol{\mu}}}-\overline{\boldsymbol{\mu}}\right\|_2^2 + 2\max_j\sqrt{\hat{w}_j}\left\|\boldsymbol{\mu}_j-\overline{\boldsymbol{\mu}}\right\|_2\left(\sqrt{\hat{w}_i}\left\|\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i\right\|_2+\sqrt{\hat{w}_i}\left\|\hat{\overline{\boldsymbol{\mu}}}-\overline{\boldsymbol{\mu}}\right\|_2\right)\right) \\ &\le 2\left\|\hat{\overline{\boldsymbol{\mu}}}-\overline{\boldsymbol{\mu}}\right\|_2^2 + \sum_{i=1}^{k}2\hat{w}_i\left\|\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i\right\|_2^2 + 2\left(\sqrt{k}\left\|\hat{\overline{\boldsymbol{\mu}}}-\overline{\boldsymbol{\mu}}\right\|_2 + \sum_{i=1}^{k}\sqrt{\hat{w}_i}\left\|\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i\right\|_2\right)\max_j\sqrt{\hat{w}_j}\left\|\boldsymbol{\mu}_j-\overline{\boldsymbol{\mu}}\right\|_2.\end{aligned}$$

The lemma follows from triangle inequality. □

The following lemma immediately follows from Lemmas 26 and 27.

Lemma 28. *Given n samples from a k-component Gaussian mixture, if Equation* (3) *and the maximum distance between two components is $\leq 10k\sigma\left(d\log\frac{n^2}{\delta}\right)^{1/4}$, then*

$$\left\|\sum_{i=1}^{k}\hat{w}_i(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})^t-\sum_{i=1}^{k}\hat{w}_i(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})^t\right\|\leq\frac{c\sigma^2dk^2\log\frac{n^2}{\delta}}{n}+c\sigma\sqrt{\frac{dk^2\log\frac{n^2}{\delta}}{n}}\max_i\sqrt{\hat{w}_i}\,\|\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}\|_2\,,$$

for a constant c.

Lemma 29. *For a set of samples $\mathbf{X}(1),\ldots\mathbf{X}(n)$ from a k-component mixture,*

$$\sum_{i=1}^{n}\frac{(\mathbf{X}(i)-\hat{\overline{\boldsymbol{\mu}}})(\mathbf{X}(i)-\hat{\overline{\boldsymbol{\mu}}})^t}{n}=\sum_{i=1}^{k}\hat{w}_i(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})^t-\hat{w}_i(\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i)(\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i)^t+\sum_{j|\mathbf{X}(j)\sim p_i}\frac{(\mathbf{X}(j)-\boldsymbol{\mu}_i)(\mathbf{X}(j)-\boldsymbol{\mu}_i)^t}{n}.$$

where $\hat{w}_i$ and $\hat{\boldsymbol{\mu}}_i$ are the empirical weights and averages of components i and $\hat{\overline{\boldsymbol{\mu}}}=\frac{1}{n}\sum_{i=1}^{n}\mathbf{X}_i$.

Proof. The given expression can be rewritten as

$$\frac{1}{n}\sum_{i=1}^{n}(\mathbf{X}(i)-\hat{\overline{\boldsymbol{\mu}}})(\mathbf{X}(i)-\hat{\overline{\boldsymbol{\mu}}})^t=\sum_{i=1}^{k}\hat{w}_i\sum_{j|\mathbf{X}(j)\sim p_i}\frac{1}{n\hat{w}_i}\mathbf{X}(j)-\hat{\overline{\boldsymbol{\mu}}})(\mathbf{X}(j)-\hat{\overline{\boldsymbol{\mu}}})^t.$$

First observe that for any set of points x_i and their average $\hat{x}$ and any value a,

$$\sum_i(x_i-a)^2=\sum_i(x_i-\hat{x})^2+(\hat{x}-a)^2.$$

Hence for samples from a component i,

$$\begin{aligned}
&\sum_{j|\mathbf{X}(j)\sim p_i}\frac{1}{n\hat{w}_i}(\mathbf{X}(j)-\hat{\overline{\boldsymbol{\mu}}})(\mathbf{X}(j)-\hat{\overline{\boldsymbol{\mu}}})^t\\
&=\sum_{j|\mathbf{X}(j)\sim p_i}\frac{1}{n\hat{w}_i}(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})^t+\sum_{j|\mathbf{X}(j)\sim p_i}\frac{1}{n\hat{w}_i}(\mathbf{X}(j)-\hat{\boldsymbol{\mu}}_i)(\mathbf{X}(j)-\hat{\boldsymbol{\mu}}_i)^t\\
&=(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})^t+\sum_{j|\mathbf{X}(j)\sim p_i}\frac{1}{n\hat{w}_i}(\mathbf{X}(j)-\hat{\boldsymbol{\mu}}_i)(\mathbf{X}(j)-\hat{\boldsymbol{\mu}}_i)^t\\
&=(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})^t+\sum_{j|\mathbf{X}(j)\sim p_i}\frac{1}{n\hat{w}_i}(\mathbf{X}(j)-\boldsymbol{\mu}_i)(\mathbf{X}(j)-\boldsymbol{\mu}_i)^t-(\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i)(\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i)^t.
\end{aligned}$$

Summing over all components results in the lemma. □

We now bound the error in estimating the eigenvalue of the covariance matrix.

Lemma 30. *Given $\mathbf{X}(1),\ldots\mathbf{X}(n)$, n samples from a k-component Gaussian mixture, if Equations* (1), (2), *and* (3) *hold, then with probability $\geq 1-2\delta$,*

$$\begin{aligned}
&\left\|\frac{1}{n}\sum_{i=1}^{n}(\mathbf{X}(i)-\hat{\overline{\boldsymbol{\mu}}})(\mathbf{X}(i)-\hat{\overline{\boldsymbol{\mu}}})^t-\hat{\sigma}^2\mathbb{I}_d-\sum_{i=1}^{k}\hat{w}_i(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})^t\right\|\\
&\leq c(n)\stackrel{\text{def}}{=}c\sigma^2\sqrt{\frac{d\log\frac{n^2}{\delta}}{n}}+c\sigma^2\frac{dk^2\log\frac{n^2}{\delta}}{n}+c\sigma\sqrt{\frac{dk^2\log\frac{n^2}{\delta}}{n}}\max_i\sqrt{\hat{w}_i}\,\|\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}\|_2\,,
\end{aligned}\tag{5}$$

for a constant c.

Proof. Since Equations (1), (2), and (3) hold, conditions in Lemmas 26 and 28 are satisfied. By Lemma 28,

$$\left\|\sum_{i=1}^{k}\hat{w}_i(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})^t-\sum_{i=1}^{k}\hat{w}_i(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})^t\right\|=\mathcal{O}\left(\sigma^2\frac{dk^2\log\frac{n^2}{\delta}}{n}+\sigma\sqrt{\frac{dk^2\log\frac{n^2}{\delta}}{n}}\max_i\sqrt{\hat{w}_i}\left\|\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}\right\|_2\right).$$

Hence it remains to show,

$$\left\|\frac{1}{n}\sum_{i=1}^{n}(\mathbf{X}(i)-\hat{\overline{\boldsymbol{\mu}}})(\mathbf{X}(i)-\hat{\overline{\boldsymbol{\mu}}})^t-\sum_{i=1}^{k}\hat{w}_i(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})^t\right\|=\mathcal{O}\left(\sqrt{\frac{kd\log\frac{5k^2}{\delta}}{n}}\sigma^2\right).$$

By Lemma 29, the covariance matrix can be rewritten as

$$\sum_{i=1}^{k}\hat{w}_i(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})(\hat{\boldsymbol{\mu}}_i-\hat{\overline{\boldsymbol{\mu}}})^t-\hat{w}_i(\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i)(\hat{\boldsymbol{\mu}}_i-\boldsymbol{\mu}_i)^t+\sum_{i=1}^{k}\sum_{j|\mathbf{X}(j)\sim p_i}\frac{1}{n}(\mathbf{X}(j)-\boldsymbol{\mu}_i)(\mathbf{X}(j)-\boldsymbol{\mu}_i)^t-\hat{\sigma}^2\mathbb{I}_d. \quad (6)$$

We now bound the norms of second and third terms in the above equation. Consider the third term, $\sum_{i=1}^{k}\sum_{j|\mathbf{X}(j)\sim p_i}\frac{1}{n}(\mathbf{X}(j)-\boldsymbol{\mu}_i)(\mathbf{X}(j)-\boldsymbol{\mu}_i)^t$. Conditioned on the fact that $\mathbf{X}(j)\sim p_i$, $\mathbf{X}(j)-\boldsymbol{\mu}_i$ is distributed $N(0,\sigma^2\mathbb{I}_d)$, therefore by Lemma 18 and Lemma 4 ,with probability $\geq 1-2\delta$,

$$\left\|\sum_{i=1}^{k}\sum_{j|\mathbf{X}(j)\sim p_i}\frac{1}{n}(\mathbf{X}(j)-\boldsymbol{\mu}_i)(\mathbf{X}(j)-\boldsymbol{\mu}_i)^t-\hat{\sigma}^2\mathbb{I}_d\right\|\leq c'\sqrt{\frac{d\log\frac{2d}{\delta}}{n}}\sigma^2+2.5\sigma^2\sqrt{\frac{\log\frac{n^2}{\delta}}{d}}.$$

The second term in Equation (6) is bounded by Lemma 25. Hence together with the fact that $d\geq 20\log n^2/\delta$ we get that with probability $\geq 1-2\delta$, the second and third terms are bounded by $\mathcal{O}\left(\sigma^2\sqrt{\frac{dk}{n}\log\frac{n^2}{\delta}}\right)$. ☐

Lemma 31. *Let* $\mathbf{u}$ *be the largest eigenvector of the sample covariance matrix and* $n\geq c\cdot dk^2\log\frac{n^2}{\delta}$. *If* $\max_i\sqrt{\hat{w}_i}\left\|\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}\right\|_2=\alpha\sigma$ *and Equation* (5) *holds, then there exists* i *such that* $|\mathbf{u}\cdot(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})|\geq\sigma(\alpha-1-1/\alpha)/\sqrt{k}$.

Proof. Observe that $\left\|\sum_j w_j\mathbf{v}_j\mathbf{v}_j^t\right\|\geq\left\|\sum_j w_j\mathbf{v}_j\mathbf{v}_j^t\frac{\mathbf{v}_i}{\|\mathbf{v}_i\|}\right\|_2\geq w_i\left\|\mathbf{v}_i\right\|_2^2$. Therefore

$$\left\|\sum_{i=1}^{k}\hat{w}_i(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})^t\right\|\geq\left\|\sum_{j=1}^{k}\hat{w}_j(\boldsymbol{\mu}_j-\overline{\boldsymbol{\mu}})(\boldsymbol{\mu}_j-\overline{\boldsymbol{\mu}})^t(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})/\left\|\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}\right\|\right\|_2\geq\alpha^2\sigma^2.$$

Hence by Lemma 30 and the triangle inequality, the largest eigenvalue of the sample-covariance matrix is $\geq\alpha^2\sigma^2-c(n)$. Similarly by applying Lemma 30 again we get,$\left\|\sum_{i=1}^{k}\hat{w}_i(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})^t\mathbf{u}\right\|_2\geq\alpha^2\sigma^2-2c(n)$. By triangle inequality and Cauchy-Schwartz inequality,

$$\begin{aligned}\left\|\sum_{i=1}^{k}\hat{w}_i(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})^t\mathbf{u}\right\|_2&\leq\sum_{i=1}^{k}\left\|\hat{w}_i(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})^t\mathbf{u}\right\|_2\\&\leq\sum_{i=1}^{k}\hat{w}_i\left\|(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})\right\|_2\max_j|(\boldsymbol{\mu}_j-\overline{\boldsymbol{\mu}})\cdot\mathbf{u}|\\&\leq\sqrt{\sum_{i=1}^{k}\hat{w}_i\left\|(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}})\right\|_2^2}\max_j|(\boldsymbol{\mu}_j-\overline{\boldsymbol{\mu}})\cdot\mathbf{u}|\\&\leq\sqrt{k}\alpha\sigma\max_j|(\boldsymbol{\mu}_j-\overline{\boldsymbol{\mu}})\cdot\mathbf{u}|.\end{aligned}$$

Hence $\sqrt{k}\alpha\sigma \max_i |(\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}) \cdot \mathbf{u}| \geq \alpha^2\sigma^2 - 2c(n)$. The lemma follows by substituting the bound on n in $c(n)$. □

We now make a simple observation on Gaussian mixtures.

Fact 32. *The samples from a subset of components A of the Gaussian mixture are distributed according to a Gaussian mixture of components A with weights being $w'_i = w_i/(\sum_{j\in A} w_j)$.*

We now prove Lemma 6.

Proof of Lemma 6. Observe that we run the recursive clustering at most n times. At every step, the underlying distribution within a cluster is a Gaussian mixture. Let Equations (1), (2) hold with probability $1 - 2\delta$. Let Equations (3) (5) all hold with probability $\geq 1 - \delta'$, where $\delta' = \delta/2n$ at each of n steps. By the union bound the total error is $\leq 2\delta + \delta' \cdot 2n \leq 3\delta$. Since Equations (1), (2) holds, the conditions of Lemmas 4 and 5 hold. Furthermore it can be shown that discarding at most $n\epsilon/4k$ samples at each step does not affect the calculations.

We first show that if $\sqrt{w_i}\,\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\|_2 \geq 25\sqrt{k^3 \log(n^3/\delta)}\sigma$, then the algorithm gets into the loop. Let w'_i be the weight of the component within the cluster and $n' \geq n\epsilon/5k$ be the number of samples in the cluster. Let $\alpha = 25\sqrt{k^3 \log(n^3/\delta)}$. By Fact 32, the components in cluster C have weight $w'_i \geq w_i$. Hence $\sqrt{w'_i}\,\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\|_2 \geq \alpha\sigma$. Since $\sqrt{w'_i}\,\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\|_2 \geq \alpha\sigma$, and by Lemma 5 $\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\| \leq 10k\sigma(d \log n^2/\delta)^{1/4}$, we have $w'_i \geq \alpha^2/(100k^2\sqrt{d \log n^2/\delta})$. Hence by lemma 24, $w'_i \geq w_i/2$ and $\sqrt{\hat{w}'_i}\,\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\|_2 \geq \alpha\sigma/\sqrt{2}$. Hence by Lemma 30 and triangle inequality the largest eigenvalue of $S(C)$ is

$$\geq \alpha^2\sigma^2/2 - c(n') \geq \alpha^2\sigma^2/4 \geq \alpha^2\hat{\sigma}^2/8 \geq 12\hat{\sigma}^2k^3 \log n^2/\delta' = 12\hat{\sigma}^2k^3 \log n^3/\delta.$$

Therefore the algorithm gets into the loop.

If $n' \geq n\epsilon/8k^2 \geq c{\cdot}dk^2 \log \frac{n^3}{\delta}$, then by Lemma 31, there exists a component i such that $|\mathbf{u}{\cdot}(\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C))| \geq \sigma(\alpha/\sqrt{2} - 1 - \sqrt{2}/\alpha)/\sqrt{k}$, where $\mathbf{u}$ is the top eigenvector of the first $n\epsilon/4k^2$ samples.

Observe that $\sum_{i\in C} w_i\mathbf{u} \cdot (\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)) = 0$ and $\max_i |\mathbf{u} \cdot (\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C))| \geq \sigma(\alpha/\sqrt{2} - 1 - \sqrt{2}/\alpha)/\sqrt{k}$. Let $\boldsymbol{\mu}_i$ be sorted according to their values of $\mathbf{u} \cdot (\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C))$, then

$$\max_i |\mathbf{u} \cdot (\boldsymbol{\mu}_i - \boldsymbol{\mu}_{i+1})| \geq \sigma \frac{\alpha/\sqrt{2} - 1 - \sqrt{2}/\alpha}{k^{3/2}} \geq 12\sigma\sqrt{\log \frac{n^3}{\delta}} \geq 9\hat{\sigma}\sqrt{\log \frac{n^3}{\delta}},$$

where the last inequality follows from Lemma 4 and the fact that $d \geq 20 \log n^2/\delta$. For a sample from component $\mathbf{p}_i$, similar to the proof of Lemma 5, by Lemma 15, with probability $\geq 1 - \delta/n^2k$,

$$\|u \cdot (\mathbf{X}(i) - \boldsymbol{\mu}_i)\| \leq \sigma\sqrt{2 \log(n^2k/\delta)}_2 \leq 2\hat{\sigma}\sqrt{\log(n^2k/\delta)},$$

where the second inequality follows from Lemma 4. Since there are two components that are far apart by $\geq 9\hat{\sigma}\sqrt{\log \frac{n^2}{\delta}}\hat{\sigma}$ and the maximum distance between a sample and its mean is $\leq 2\hat{\sigma}\sqrt{\log(n^2k/\delta)}$ and the algorithm divides into at-least two non-empty clusters such that no two samples from the same distribution are clustered into two clusters.

For the second part observe that by the above concentration on $\mathbf{u}$, no two samples from the same component are clustered differently irrespective of the mean separation. Note that we are using the fact that each sample is clustered at most $2k$ times to get the bound on the error probability. The total error probability by the union bound is $\leq 4\delta$. □

D.4 Proof of Lemma 7

We show that if the conclusions in Lemmas 6 and 24 holds, then the lemma is satisfied. We also assume that the conclusions in Lemma 30 holds for all the clusters with error probability $\delta' = \delta/k$. By the union bound the total error probability is $\leq 7\delta$.

By Lemma 6 all the components within each cluster satisfy $\sqrt{w_i}\left\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\right\|_2 \leq 25\sigma\sqrt{k^3\log(n^3/\delta)}$. Let $n \geq c{\cdot}dk^9\epsilon^{-4}\log^2 d/\delta$. For notational convenience let $S(C) = \frac{1}{|C|}\sum_{i=1}^{|C|}(\mathbf{X}(i)-\overline{\boldsymbol{\mu}}(C))(\mathbf{X}(i)-\overline{\boldsymbol{\mu}}(C))^t - \hat{\sigma}^2\mathbb{I}_d$. Therefore by Lemma 30 for large enough c,

$$\left\|S(C) - \frac{n}{|C|}\sum_{i\in C}\hat{w}_i(\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C))(\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C))^t\right\| \leq \frac{\epsilon^2\sigma^2}{1000k^2}\frac{n}{|C|}.$$

Let $\mathbf{v}_1, \mathbf{v}_2, \ldots \mathbf{v}_{k-1}$ be the top eigenvectors of $\frac{1}{|C|}\sum_{i\in C} w_i(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}(C))(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}(C))^t$. Let $\eta_i = \sqrt{\hat{w}_i'}\left\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\right\|_2 = \sqrt{\hat{w}_i}\sqrt{\frac{n}{|C|}}\left\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\right\|_2$. Let $\boldsymbol{\Delta}_i = \frac{\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}(C))}{\|(\boldsymbol{\mu}_i-\overline{\boldsymbol{\mu}}(C))\|_2}$. Therefore,

$$\sum_{i\in C}\frac{n}{|C|}\sum_{i\in C}\hat{w}_i(\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C))(\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C))^t = \sum_{i\in C}\eta_i^2\boldsymbol{\Delta}_i\boldsymbol{\Delta}_i^t.$$

Hence by Lemma 20, the projection of $\boldsymbol{\Delta}_i$ on the space orthogonal to top $k-1$ eigenvectors of $S(C)$ is

$$\leq \sqrt{\frac{\epsilon^2\sigma^2}{1000k^2}\frac{n}{|C|}}\frac{1}{\eta_i} \leq \frac{\epsilon\sigma}{16\sqrt{\hat{w}_i}\left\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\right\|_2 k} \leq \frac{\epsilon\sigma}{8\sqrt{2}\sqrt{w_i}\left\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\right\|_2 k}.$$

The last inequality follows from the bound on $\hat{w}_i$ in Lemma 24.

D.5 Proof of Theorem 8

We show that the theorem holds if the conclusions in Lemmas 7 and 26 holds with error probability $\delta' = \delta/k$. Since in the proof of Lemma 7, the probability that Lemma 6 holds is included, Lemma 6 also holds with the same probability. Since there are at most k clusters, by the union bound the total error probability is $\leq 9\delta$.

For every component i, we show that there is a choice of mean vector and weight in the search step such that $w_i D(\mathbf{p}_i, \hat{\mathbf{p}}_i) \leq \epsilon/2k$ and $|w_i - \hat{w}_i| \leq \epsilon/4k$. That would imply that there is a $\hat{\mathbf{f}}$ during the search such that

$$D(\mathbf{f}, \hat{\mathbf{f}}) \leq \sum_C\sum_{i\in C} w_i D(\mathbf{p}_i, \hat{\mathbf{p}}_i) + 2\sum_{i=1}^{k-1}|w_i - \hat{w}_i| \leq \frac{\epsilon}{2k} + \frac{\epsilon}{2k} = \epsilon.$$

Since the weights are gridded by $\epsilon/4k$, there exists a $\hat{w}_i$ such that $|w_i - \hat{w}_i| \leq \epsilon/4k$. We now show that there exists a choice of mean vector such that $w_i D(\mathbf{p}_i, \hat{\mathbf{p}}_i) \leq \epsilon/2k$. Note that if a component has weight $\leq \epsilon/4k$, the above inequality follows immediately. Therefore we only look at those components with $w_i \geq \epsilon/4k$, by Lemma 24, for such components $\hat{w}_i \geq \epsilon/5k$ and therefore we only look at clusters such that $|C| \geq n\epsilon/5k$. By Lemmas 14 and for any i,

$$D(\mathbf{p}_i, \hat{\mathbf{p}}_i)^2 \leq 2\sum_{j=1}^{d}\frac{(\mu_{i,j} - \hat{\mu}_{i,j})^2}{\sigma^2} + 8d\frac{(\sigma^2 - \hat{\sigma}^2)^2}{\sigma^4}.$$

Note that since we are discarding at most $n\epsilon/8k^2$ random samples at each step. A total number of $\leq n\epsilon/8k$ random samples are discarded. It can be shown that this does not affect our calculations and we ignore it in

this proof. By Lemma 4, the first estimate of σ^2 satisfies $|\hat{\sigma}^2 - \sigma^2| \leq 2.5\sigma^2\sqrt{\log n^2/\delta}$. Hence while searching over values of $\hat{\sigma}^2$, there exist one such that $|\sigma'^2 - \sigma^2| \leq \epsilon\sigma^2/\sqrt{64dk^2}$. Hence,

$$D(\mathbf{p}_i, \hat{\mathbf{p}}_i)^2 \leq 2\frac{\|\boldsymbol{\mu}_i - \hat{\boldsymbol{\mu}}_i\|_2^2}{\sigma^2} + \frac{\epsilon^2}{8k^2}.$$

Therefore if we show that there is a mean vector $\hat{\boldsymbol{\mu}}_i$ during the search such that $\|\boldsymbol{\mu}_i - \hat{\boldsymbol{\mu}}_i\|_2 \leq \epsilon\sigma/\sqrt{16k^2\hat{w}_i}$, that would prove the Lemma. By triangle inequality,

$$\|\boldsymbol{\mu}_i - \hat{\boldsymbol{\mu}}_i\|_2 \leq \left\|\overline{\boldsymbol{\mu}}(C) - \hat{\overline{\boldsymbol{\mu}}}(C)\right\|_2 + \left\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C) - (\hat{\boldsymbol{\mu}}_i - \hat{\overline{\boldsymbol{\mu}}}(C))\right\|_2.$$

By Lemma 26 for large enough n,

$$\left\|\overline{\boldsymbol{\mu}}(C) - \hat{\overline{\boldsymbol{\mu}}}(C)\right\|_2 \leq c\sigma\sqrt{\frac{dk\log^2 n^2/\delta}{|C|}} \leq \frac{\epsilon\sigma}{8k\sqrt{w_i}}.$$

The second inequality follows from the bound on n and the fact that $|C| \geq n\hat{w}_i$. Since $w_i \geq \epsilon/4k$, by Lemma 24, $\hat{w}_i \geq w_i/2$, we have

$$\|\boldsymbol{\mu}_i - \hat{\boldsymbol{\mu}}_i\|_2 \leq \left\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C) - (\hat{\boldsymbol{\mu}}_i - \hat{\overline{\boldsymbol{\mu}}}(C))\right\|_2 + \frac{\epsilon\sigma}{8k\sqrt{w_i}}.$$

Let $\mathbf{u}_1 \ldots \mathbf{u}_{k-1}$ are the top eigenvectors the sample covariance matrix of cluster C. We now prove that during the search, there is a vector of the form $\sum_{j=1}^{k-1} g_j\epsilon_g\hat{\sigma}\mathbf{u}_j$ such that $\left\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C) - \sum_{j=1}^{k-1} g_j\epsilon_g\hat{\sigma}\mathbf{u}_j\right\|_2 \leq \frac{\epsilon\sigma}{8k\sqrt{w_i}}$, during the search, thus proving the lemma. Let $\eta_i = \sqrt{w_i}\,\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\|_2$. By Lemma 7, there are set of coefficients α_i such that

$$\frac{\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)}{\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\|_2} = \sum_{j=1}^{k-1}\alpha_j\mathbf{u}_j + \sqrt{1 - \|\alpha\|^2}\mathbf{u}',$$

where $\mathbf{u}'$ is perpendicular to $\mathbf{u}_1 \ldots \mathbf{u}_{k-1}$ and $\sqrt{1 - \|\alpha\|^2} \leq \epsilon\sigma/(8\sqrt{2}\eta_i k)$. Hence, we have

$$\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C) = \sum_{j=1}^{k-1}\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\|_2\,\alpha_j\mathbf{u}_j + \|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\|_2\sqrt{1 - \|\alpha\|_2^2}\mathbf{u}',$$

Since $w_i \geq \epsilon/4k$ and by Lemma 6, $\eta_i \leq 25\sqrt{k^3}\sigma\log(n^3/\delta)$, and $\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\|_2 \leq 100\sqrt{k^4\epsilon^{-1}}\sigma\log(n^3/\delta)$. Therefore $\exists g_j$ such that $|g_j\hat{\sigma} - \alpha_j| \leq \epsilon_g\hat{\sigma}$ on each eigenvector. Hence,

$$\begin{aligned}
w_i\left\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C) - \sum_{i=1}^{k-1} g_j\epsilon_g\hat{\sigma}\mathbf{u}_j\right\|_2^2 &\leq w_i k\epsilon_g^2\hat{\sigma}^2 + w_i\|\boldsymbol{\mu}_i - \overline{\boldsymbol{\mu}}(C)\|_2^2(1 - \|\alpha\|^2) \\
&\leq k\epsilon_g^2\hat{\sigma}^2 + \eta_i^2\frac{\epsilon^2\sigma^2}{128\eta_i^2k^2} \\
&\leq \frac{\epsilon^2\sigma^2}{128k^2} + \frac{\epsilon^2\sigma^2}{128k^2} \leq \frac{\epsilon^2\sigma^2}{64k^2}.
\end{aligned}$$

The last inequality follows by Lemma 4 and the fact that $\epsilon_g \leq \epsilon/16k^{3/2}$, and hence the theorem. The run time can be easily computed by retracing the steps of the algorithm and using an efficient implementation of single-linkage.